arXiv:1610.05604v1 [stat.ML] 18 Oct 2016

Dynamic Assortment Personalization in High Dimensions

Nathan Kallus
Cornell University and Cornell Tech
kallus@cornell.edu

Madeleine Udell
Cornell University
udell@cornell.edu

We demonstrate the importance of structural priors for effective, efficient large-scale dynamic assortment personalization. Assortment personalization is the problem of choosing, for each individual or consumer segment (type), a best assortment of products, ads, or other offerings (items) so as to maximize revenue. This problem is central to revenue management in e-commerce, online advertising, and multi-location brick-and-mortar retail, where both items and types can number in the thousands-to-millions. Data efficiency is paramount in this large-scale setting. A good personalization strategy must dynamically balance the need to learn consumer preferences and to maximize revenue.

We formulate the dynamic assortment personalization problem as a discrete-contextual bandit with m contexts (customer types) and many arms (assortments of the n items). We assume that each type's preferences follow a simple parametric model with n parameters. In all, there are mn parameters, and existing literature suggests that order optimal regret scales as mn. However, this figure is orders of magnitude larger than the data available in large-scale applications, and imposes unacceptably high regret.

In this paper, we impose natural structure on the problem – a small latent dimension, or low rank. In the static setting, we show that this model can be efficiently learned from surprisingly few interactions, using a time- and memory-efficient optimization algorithm that converges globally whenever the model is learnable. In the dynamic setting, we show that structure-aware dynamic assortment personalization can have regret that is an order of magnitude smaller than structure-ignorant approaches. We validate our theoretical results empirically.

Key words: Personalization, Contextual bandit, Assortment planning, Discrete choice, High-dimensional learning, Large-scale learning, First-order optimization, Recommender Systems, Matrix completion

1. Introduction

In many commerce, e-commerce, and advertising settings, customers or users are presented with an assortment (set) of products, ads, or other offerings. Customers choose which to products to buy, ads to click, or (generically) items to interact with, from among the assortment that is presented.

Firms choose which assortment to present to the customer, and collect the revenue (or other benefit or loss) resulting from each customer's choice. Choosing the assortment that maximizes expected revenue is a central problem in revenue management. This problem goes by the name *assortment planning* or *assortment optimization*. When assortments are tailored to each individual or consumer segment, the problem is known as *assortment personalization*. Successful personalization hinges on learning the preferences of each customer over the set of all offerings.

Assortment personalization requires a good understanding of how consumer tastes vary. These preferences are learned from interactions with consumers. Usually, the number of interactions between the firm and customer is limited, so efficient estimation of customer preferences is critical. But estimating customer preferences is no easy task: there are combinatorially many assortments of items, and so without further assumptions, combinatorially many quantities to estimate. To enable tractable estimation, customer preferences are generally modeled parametrically, often using the multinomial logit (MNL) model. This model posits a single nominal vector of preferences which governs the choices of all consumers. Individual differences are modeled as random, homogeneous deviations from these universal preferences, and treated as noise. These one-size-fits-all models offer no opportunity for personalization and fit heterogeneous populations poorly.

Learning a personalized model can improve performance significantly. A personalized model *segments* the population into *types* (which may be single individuals, or geographic, demographic, or temporal groups), and fits a separate model for each type. A type is usually chosen to be the smallest group that is uniquely identifiable given the available data. In the offline brick-and-mortar retail setting, populations may be segmented by store branch, since firms have data on aggregate choices in each store branch, and tastes often vary geographically and by setting (mall, street, *etc.*). Customers may also be segmented using loyalty program data. In the e-commerce and online advertising settings, firms have data on the choice of each individual customer from individually personalized assortments. Hence in the limit, each type may represent a single customer.

However, the number of interactions needed to estimate all of these models separately is linear in the product of the number of types and the number of items. This is an astronomical figure in most e-commerce and online advertising contexts, where types and items both number in the millions. It can be enormous even in brick-and-mortar settings: large chains may comprise hundreds or thousands of locations and offer hundreds of thousands of stock keeping units (SKUs). Worse, these numbers are far greater than the number of interactions. Each user can view only so many webpages or consider only so many products — generally, far fewer than the number available. The data from each of these interactions is also limited to a single solitary choice (or lack thereof) out of the assortment.

In this paper, we propose a new approach to assortment personalization. To enable tractable estimation of a personalized model with limited data, we proposal a new structural model. In this model, the choices of each type are governed by its own personalized preference vector (with one dimension for every item); but these preference vectors span (or lie close to) only a low-dimensional subspace. We demonstrate how to estimate the parameters of this model with computationally tractable algorithms, and provide a proof of recovery with high-probability from few samples (sub-linear in number type-item combinations). Numerically, we show that given the same data, our estimator performs much better than either a single MNL model, or decoupled MNL models for each type, which we call "structure-ignorant" methods.

We also consider a *dynamic* assortment personalization problem: starting with *no* data, how should we choose assortments to offer to different types to maximize total profit, or equivalently, to minimize total regret? We show theoretically and numerically that our algorithm achieves regret orders-of-magnitude smaller than structure-ignorant methods. For example, when the parameter matrix has rank r, we achieve regret of order $r\max(m,n)\log T$ after T interactions, where ignoring structure would have yielded regret of order $\min(m,n)\max(m,n)\log T$.[1] Our results demonstrate that assortment personalization enables orders of magnitude better performance than competing approaches, and can be achieved with a tractable, efficient estimator.

All proofs are given in the Appendix.

1.1. The Main Problem

In this section we describe the problem of *dynamic assortment personalization*. We consider m types, n items, T interactions, and assortment capacity of K. Presented an assortment, a consumer can choose any of the items presented or choose nothing, which we refer to as choosing the 0^{th} item. The revenue generated when a consumer of type i chooses item j is W_{ij} and is 0 if no item is chosen. At the beginning, only m, n, T, and $W\in\mathbb{R}^{m\times n}$ are revealed. The problem instance is also described by two additional *unknown* parameters $\mu^\star\in\Delta^m=\left\{\alpha\in\mathbb{R}^m_+:\sum_{i=1}^m\alpha_i=1\right\}$ and $\Theta^\star\in\mathbb{R}^{m\times n}$. The parameter $\mu^\star$ describes the arrival distribution over types, and $\Theta^\star$ is the matrix of the parameters governing the MNL choice model of each type over all items.

The problem proceeds as follows. At each $t=1,\dots,T$:

1. a type i_t is drawn at random from $\{1,\dots,m\}$ with probability proportional to weights $\mu^\star_i$ and i_t is reveled,

2. we choose any $S_t\subset\{1,\dots,n\}$ with $|S_t|\le K$,

[1] Note that for these to make sense as orders, we let m,n,T all vary simultaneously.

3. an item j_t is drawn at random from $\{1,\dots,n\}$ with probability proportional to

$$\text{weight}(j) = \begin{cases} 1 & j = 0 \\ 0 & j \neq 0,\, j \notin S_t \\ \exp(\Theta^\star_{i_t j}) & j \neq 0,\, j \in S_t \end{cases}$$

and j_t is revealed, and finally

4. if $j_t = 0$ we get reward $r_t = 0$ and otherwise we get reward $r_t = W_{i_t j_t}$.

Define a set of random variables U_t for $t = 1,\dots,T$ independent of all variables above, and define the filtration

$$\mathcal{F}_t = \sigma(U_t, i_t, i_{t-1}, j_{t-1}, S_{t-1}, \dots, i_1, j_1, S_1),$$

the smallest σ algebra with respect to which all variables known at time t, but before the choice of set S_t is made, are measurable. An algorithm π is an assignment to the random variables S_t so that, for every t, S_t is measurable with respect to $\mathcal{F}_t$. (Notice that the random variables U_t allow randomized algorithms.) Given a problem instance and an algorithm π, define $\mathbb{P}^\pi$ and $\mathbb{E}^\pi$ as the probability and expectation when the sets S_t are selected according to algorithm π.

Define the Θ-greedy algorithm, π_Θ, which chooses

$$S_t \in \operatorname*{argmax}_{|S| \leq K} \sum_{j \in S} \frac{e^{\Theta_{i_t j}} W_{i_t j}}{1 + \sum_{j' \in S} e^{\Theta_{i_t j'}}}.$$

The Θ-greedy algorithm π_Θ always chooses the assortment S_t that would maximize revenue if Θ were the true choice parameter matrix.

Definition 1. Given an instance $(m, n, T, W, \mu^\star, \Theta^\star)$, the *regret* of the algorithm π is

$$\text{Regret}(\pi) = \mathbb{E}^{\pi\Theta^\star}\left[\sum_{t=1}^T r_t\right] - \mathbb{E}^\pi\left[\sum_{t=1}^T r_t\right]$$

One of our main results will be to construct an algorithm $\pi_{\text{nuc-norm}}$ that exploits low-rank structure in $\Theta^\star$ to achieve order-lower regret compared with structure-ignorant algorithms.

Theorem 1 **(Informal)**. *If* $\text{rank}(\Theta^\star) \leq r$ *and under some additional conditions,*

$$\text{Regret}(\pi_{nuc\text{-}norm}) = O(r \max(m,n) \log T).$$

Here we see that the regret grows sublinearly in the dimension mn of the problem, for fixed rank. Contrast this result with the best rate achievable by an algorithm that ignores the low rank structure of the problem. This second theorem uses a simple argument to extend Theorem 1 of Sauré and Zeevi (2013) to a setting with many types.

Theorem 2 **(Informal)**. *If* π *is a structure-ignorant algorithm and under some additional conditions,*

$$\text{Regret}(\pi) = \Omega(\min(m,n) \max(m,n) \log T).$$

Turn to Theorems 5 and 6 in Section 3 for the formal statement, definition of structure-ignorant, and the additional conditions under which the theorems hold.

1.2. Contributions

Let us summarize the contributions of this paper. We propose the *low-rank mixed multinomial logit* (LRMMNL) model for discrete choice in a heterogeneous population. We formulate the *high-dimensional dynamic assortment personalization problem* as a discrete-contextual bandit where choice is governed by the LRMMNL model. We propose a nuclear-norm regularized maximum likelihood estimator for the static estimation problem of learning the LRMMNL model from only choice observations. We prove Theorem 3 to bound the sample complexity of learning this model from choice data and show that learning can be done using sublinearly many observations. We develop Algorithm 1 for efficient computation of our estimator from very large-scale data. We develop Algorithm 2 for the dynamic assortment personalization problem and prove that it achieves regret that is sublinear in problem size and logarithmic in time horizon. We study the regret of the dynamic assortment personalization problem and compare structure-aware and structure-ignorant algorithms to show an order-of-magnitude discrepancy in regret. We also study the regret of the non-contextual dynamic assortment planning problem where choice is governed by the LRMMNL model, and develop algorithms with low regret. We study both the static estimation problem and the dynamic assortment personalization problem numerically and demonstrate the significant advantage offered by the leverage of structure for dynamic assortment personalization.

1.3. Related Work

More basic than the question of how to optimize assortments is the question of how customers choose from among a set of goods. Discrete choice models posit answers to this question in the form of a probability distribution over choices. Luce (1959) proposed an early discrete choice model based on an axiomatic theory, resulting in the basic attraction model. Later, the work of McFadden (1973) on random utility theory led to the MNL model, which posits that customer choices follow a log-linear model in a vector of customer preference parameters. Fitting a single MNL model is as simple as counting the number of times an item is chosen relative to the other offerings. (These counts give the maximum likelihood estimate for the model.)

The MNL model has many variants. The mixture of MNLs (MMNL) model models consumer choice as a mixture of MNL models with different parameters. With sufficiently many mixture components, an MMNL model can approximate arbitrarily closely any choice model that arises from a distribution over individual preferences (McFadden and Train 2000, Farias et al. 2013, van Ryzin and Vulcano 2014). Other derivatives of the MNL model include the nested logit model (Williams 1977) and its extensions (McFadden 1980). Recently a new choice model was proposed that arises when substitutions from one good to another are assumed to form a markov chain Blanchet et al. (2013).

Choosing the optimal assortment can be computational hard or computationally easy depending on the choice model. Under the MNL model, it is easy to optimize assortments: Talluri and Van Ryzin (2006) show that presenting items in revenue sorted order is always optimal. On the other hand, it is NP-hard to optimize a single assortment to be offered to one MMNL population, even with only two mixture components, but approximation schemes exist (Rusmevichientong et al. 2014). Assortment optimization over the nested MNL model is computationally hard in general (Davis et al. 2014) but easy in some cases (Li et al. 2015). Optimizing an assortment of constrained cardinality under the MNL model is easy (Megiddo 1979, Rusmevichientong et al. 2010), while optimizing an assortment with weighted budget constraint is hard (Désir and Goyal 2014).

Assortment optimization with limited inventory is even more complex. Such problems need to be solved over multiple periods, a different assortment offered each period to take into account possible future stock outs. Talluri and Van Ryzin (2006) solve the classic problem with a single MNL model, while Bernstein et al. (2011), Golrezaei et al. (2014) solve problem with a mixture model representing heterogeneous populations. In all of these, all preferences of all populations are assumed known.

When preferences are unknown and are to be learned simultaneously with assortment optimization, we can conceive of assortments as arms and consumer choice as bandit feedback to get the problem of *dynamic assortment optimization* (without context). Rusmevichientong et al. (2010) first formulated this problem of assortment optimization while learning discrete choice for a single MNL model and showed that their algorithm has regret upper bounded in order by $n \log^2 T$ for n items and T interactions. Sauré and Zeevi (2013) improved this to $n \log T$ and showed that this order is optimal. Caro and Gallien (2007) were the first to consider dynamic optimization of assortments under learning by studying a problem related to that of Rusmevichientong et al. (2010), Sauré and Zeevi (2013). Their problem differs in that the demand to be learned is assumed exogenous and independent of the combination offered and other items. The assortments constitute a simultaneous play of multiple arms, rather than an optimal assortment from which a customer chooses zero or one items.

Our *dynamic assortment personalization* is an instance of a *contextual* bandit problem with discrete contexts. It arises when at each interaction, a different context, drawn from some finite set, is observed. Based on this discrete contextual information, the problem is to personalize the assortment to target each context as well as possible, while also learning to improve performance. In particular, a good algorithm will use lessons learned in one context to improve performance in other contexts. To the best of our knowledge, we are the first to consider this dynamic assortment personalization problem, and in particular the first to consider any stochastic bandit with discrete contexts, rather than continuous contexts with a functional relationship to rewards (such as linear).

Lai and Robbins (1985) posed the classic stochastic multi-armed bandit problem, in which each of n arms has an initially-unknown bounded reward distribution and in each of T time steps one chooses one arm to pull with the overall goal of minimal regret: the expected difference in reward between the prescient policy that always pulls the best arm and one's actual performance. An alternative formulation of the multi-armed bandit problem involves rewards that, instead of being distributed according to a fixed unknown distribution, may change adversarially in response to the choice of arm. For a discussion of the differences between stochastic and adversarial bandits we refer the reader to Bubeck and Cesa-Bianchi (2012). In this paper we focus solely on the *stochastic* bandit. Examples of contextual stochastic bandits include Rigollet and Zeevi (2010), Perchet and Rigollet (2013), Goldenshluger and Zeevi (2013), Slivkins (2014), Bastani and Bayati (2015), which all focus on the setting with n generally unrelated arms, where each arm is associated with a regression function that governs the expected reward conditioned on a continuous vector of covariates representing context. The former two papers assume a general non-parametric functional dependence; the latter three assume a linear regression function. In all these papers, the context is parametrized by a continuous (scalar or vector) quantity; in other words, the relation between different contexts is embedded topologically and known in advance.

In contrast, in the dynamic assortment personalization, the relation between different contexts must be learned from the data. Contexts are discrete; they correspond to rows of an unknown parameter matrix which governs consumer choice. The observation of choice can be likened (imperfectly) to the noisy observation of an entry of the matrix. This analogy brings to mind the problem of matrix completion: the problem of (approximately) recovering an (approximately) low rank matrix from a few (noisy) samples from its values.

Some of our results are in the same vein as statistical matrix completion bounds Following groundbreaking work on exact completion of exactly low rank matrices whose entries are observed without noise (Candès and Tao 2010, Candès and Recht 2009, Recht et al. 2010, Keshavan et al. 2010), approximate recovery results have been obtained for a variety of different noisy observation models. These include observations with additive gaussian (Candès and Plan 2009) and subgaussian (Keshavan et al. 2009a) noise, 0-1 (Bernoulli) observations (Davenport et al. 2014), observations from any exponential family distribution (Gunasekar et al. 2014), and observations generated according to the Bradley-Terry-Luce model for pairwise comparisons (Lu and Negahban 2014, Oh et al. 2015). The latter is most related to our matrix recovery results, but is dependent on the observation of the full rankings of all items in any given list, does not allow the customer to choose to buy nothing, and requires the possibility of duplicate items (with nonzero probability). In contrast, we consider observing only the choice or lack thereof from assortments that are actual subsets of the items and include the option of "no buy." This type of data is much more common

in practice because it corresponds to passive observations of consumer behavior that is truthful (assuming choice is utility maximizing) and it is easy to collect because it can be derived from transactional data.

2. The Low-Rank Mixed Multinomial Logit Choice Model

In this section we describe the LRMMNL model, study the static estimation problem under observing only choice, propose an estimator, prove recovery bounds, and develop a fast algorithm for computing the estimator from large-scale data.

The mixed MNL (MMNL) model over types $i=1,\dots,m$ and items $j=1,\dots,n$ is parameterized by $\mu^\star \in \Delta^m$ and $\Theta^\star \in \mathbb{R}^{m\times n}$ and describes two random variables: type I and choice J. For any given assortment $S \subset \{1,\dots,n\}$, choice J is assumed to be distributed according to the following model

$$\begin{aligned} \mathbb{P}(J=j;S) &= \sum_{i=1}^n \mathbb{P}(I=i)\,\mathbb{P}(J=j \mid I=i;S) \\ \mathbb{P}(J=j \mid I=i;S) &= \frac{1}{\sum_{j'\in S} e^{-\Theta^\star_{ij'}}} \times \begin{cases} 1 & j=0 \\ 0 & j\neq 0,\, j\notin S_t \\ \exp(\Theta^\star_{i_t j}) & j\neq 0,\, j\in S_t \end{cases} \\ \mathbb{P}(I=i) &= \mu^\star_i \end{aligned} \tag{1}$$

where $J=0$ represents the choice not to choose from S – an option that is always available for any assortment S.

The LRMMNL model posits an MMNL model in which the parameter $\Theta^\star$ has low rank:

$$\operatorname{rank}(\Theta^\star) \ll m, n.$$

We will also consider the case where $\Theta^\star$ has approximately low rank:

$$\overline{\sigma}_{r+1}(\Theta^\star) \approx 0,$$

where $\overline{\sigma}_{r+1}(\Theta^\star)$ is the sum of the singular values of $\Theta^\star$ smaller than the r^{th} largest singular value, i.e.,

$$\overline{\sigma}_{r+1}(\Theta^\star) = \sum_{j=r+1}^{\min\{m,n\}} \sigma_j(\Theta^\star). \tag{2}$$

2.1. Implications of the LRMMNL Model

As per the approximation results of McFadden and Train (2000), for any choice distribution over a population, there is a variable I such that the choice distribution is approximately MNL conditioned on I. Hence, if we segment the population finely enough, the MMNL model above is as general as need be.

The MMNL model can be interpreted as a random utility maximizing model in the following way. Defining $u_{ij} = \Theta^\star_{ij}$ as the mean utility derived from item j for the type i, we can describe the choice of a random customer of type i as maximizing her own utility among the items on offer, where her own utility is the type's mean utility plus random idiosyncrasies distributed per the extreme value distribution:

$$J = \max_{j \in S_t} (u_{Ij} + \zeta_j) \quad \text{where} \quad \zeta_j \sim \text{Gumbell}(0,1). \tag{3}$$

Our model (1) can therefore be thought of in either of two ways. One is as assuming that customers are clustered into types within each of which customers have a private, idiosyncratic utility distributed as in (3) and the heterogeneity of the population is described by the varying mean utilities u_{ij} over types i. Another interpretation is that each customer is her own type. The random idiosyncrasies associated with each choice event with the same customer reflects the oft-observed human inconsistencies in decision making or slight variations over time in preferences (Kahneman and Tversky 1979, DeShazo and Fermo 2002) – it would be unrealistic to assume that in each interaction with the same customer, we observe choice exactly without noise from a completely fixed preference list.

Learning the preferences of multiple, heterogeneous customer types simultaneously can be very hard without additional structure. Both Bernstein et al. (2011) and Golrezaei et al. (2014) study multi-period assortment optimization problems with multiple, heterogeneous customer types under the assumption of full knowledge of the distributions of consumer choice. Both consider case studies in which they estimate these distributions from static data in order to evaluate the performance of their approaches on distributions that mimic real data. However, in both cases, this is limited to very few segments (3 and 10, respectively). One reason that this may be is that the increasing dimension of the choice model with many segments or individual consumers make estimation practically infeasible for models that are much larger. We address this by allowing very many types and items, while requiring that the underlying dimension of the model is low in the sense that our parameter matrix has (approximately) low rank.

2.2. The Static Estimation Problem

Next, we describe an observation model and the problem of estimating the LRMMNL parameters from observed data. We suppose that we have N observations $\{(i_t, j_t, S_t) : t = 1, \ldots, N\}$ where S_t is sampled uniformly at random from the set of subsets of $\{1, \ldots, n\}$ of size K_t, the sequence K_t is arbitrary (possibly random) satisfying $K_t \leq K$, and i_t, j_t are iid (independent and identically distributed) according to the model (1). We also assume that $||\Theta^\star||_\infty \leq \alpha/\sqrt{mn}$ for purely technical reasons. The assumption of (normalized) bounded entries assumption is standard in many matrix

completion recovery results (see Section 1.3) and is necessary for our proof of recovery with high probability; see below.

It is important to highlight that our observation model consists of observing only the choice made by customers. In practical applications, this is typically the only observation possible. Moreover, it is generally truthful since it is utility maximizing, unlike reporting rankings in a survey or focus group.

2.2.1. Our Estimator Define the negative log likelihood of the observations given parameter Θ as

$$L(\Theta) = \frac{1}{N}\sum_{t=1}^{N} \log\left(\left(1+\sum_{j\in S_t} e^{\Theta_{ij}}\right)\left(\begin{cases} 1 & j_t = 0 \\ e^{\Theta_{ij_t}} & \text{otherwise} \end{cases}\right)^{-1}\right). \tag{4}$$

We define our estimator $\widehat{\Theta}$ for $\Theta^\star$ as any solution of the nuclear norm regularized maximum likelihood problem

$$\begin{array}{l} \text{minimize } L(\Theta) + \lambda\|\Theta\|_*, \\ \text{subject to } \|\Theta\|_\infty \le \alpha/\sqrt{mn}, \end{array} \tag{5}$$

where $\lambda > 0$ is a tuning parameter and the nuclear norm $\|\Theta\|_*$ is the sum of the singular values of Θ. The constraint $\|\Theta\|_\infty \le \alpha/\sqrt{mn}$ appears purely as an artifact of the proof; we recommend to omit this constraint in practice. We omit this constraint both in our specialized algorithm (Section 2.3) and in our numerical results (Section 5), which demonstrates the practical irrelevance of this constraint.

Problem (5) is convex and hence can be solved by a variety of standard convex methods that take advantage of the special structure of the problem (Cai et al. 2010, Parikh and Boyd 2014, Hazan 2008, Orabona et al. 2012). In Section 2.3 we provide a specialized first-order algorithm that, in fact, works on the *non-convex*, factored form of the problem for increased speed, but still guarantees convergence to the global optimum with high probability.

Our estimator $\hat{\mu}$ for the customer type distribution $\mu^\star$ is the empirical frequencies of each type:

$$\hat{\mu}_i = \frac{1}{N}\sum_{t=1}^{N} \mathbb{I}\left[i_t = i\right].$$

2.2.2. Main Results for the Static Estimation Problem In this section, we bound the error of our estimators $\widehat{\Theta}$ and $\hat{\mu}$. Our first bound depends on the following quantities, which capture the complexity of learning the preferences of all customer types over all items.

1. *Number of observations.* The bound decreases as the number N of observations increases.

2. *Number of parameters.* The bound grows as the dimensions m, n of the parameter matrix $\Theta^\star$ grow.

3. *Underlying rank dimension.* For any $r \leq \min(m,n)$, our bound decomposes into two error terms. The first error term is the error in estimating the top r "principal components" of the parameter matrix. This error grows with $\sqrt{r}$ and captures the benefit of learning only the most salient features instead of all parameters at once. The second error is the error in approximating the parameter matrix by only its top r "principal components." In particular, if it is assumed that $\Theta^\star$ is exactly rank r, then this latter error is zero. More generally, however, we may be interested in estimating parameter matrices that are only approximately low rank, *i.e.*, with quickly decaying singular values past the top r, in which case our bound depends on the sum of the remaining singular values.

4. *Size of parameters.* Our bound grows with the (scaled) maximum magnitude of any entry α.

5. *Size of assortments.* Our bound grows with the maximum size K of the assortments.

THEOREM 3. *Let $\tau \geq 1$ be given, $\rho \geq 1$ be such that $1/\rho \leq m\mu_i \leq \rho \ \forall i = 1,\dots,m$, and α be such that $||\Theta^\star||_\infty \leq \alpha/\sqrt{mn}$. Fix $\lambda = 8\sqrt{\frac{\tau\rho K(m+n)\log(m+n)}{mnN}}$. Suppose $N \leq mn\log(m+n)$. Then under the observation model in Section 2.2 and for any integer $r \leq \min\{m,n\}$, with probability at least $1-3(m+n)^{-\tau}$, any solution $\widehat{\Theta}$ to Problem (5) satisfies*

$$\|\widehat{\Theta} - \Theta^\star\|_F \leq 2048\sqrt{\tau}\alpha e^{\frac{4\alpha}{\sqrt{mn}}} \max\left\{ \sqrt{\frac{r(m+n)\log(m+n)}{N}(\rho K)^3}, \left(\frac{\overline{\sigma}_{r+1}(\Theta^\star)(m+n)\log(m+n)}{N}(\rho K)^3\right)^{1/4} \right\},$$

where $\overline{\sigma}_{r+1}(\Theta^\star)$ is the sum of the remaining singular values of $\Theta^\star$ smaller than the r^{th} largest singular value, as defined in eq. (2).

Note that if $\text{rank}(\Theta^\star) \leq r$, then $\overline{\sigma}_{r+1}(\Theta^\star) = 0$ and the second term drops. So when $\text{rank}(\Theta^\star) \leq r$, Theorem 3 shows that we need only have $N = \omega(r\max(m,n)\log(m+n))$ observations in order to have a consistent estimate for $\Theta^\star$.

A closely related result to Theorem 3 appeared in our preliminary work (Kallus and Udell 2016), which focuses only on the static estimation problem and only in the absence of the no-choice option. In assortment personalization, we must consider estimation in the permanent presence of a no-choice option in *any* assortment and where the mixtures μ are not necessarily uniform.

Our second bound provides recovery guarantees for the customer type distribution $\mu^\star$.

THEOREM 4. *Let $\tau \geq 0$ and $q \in [1,\infty]$ be given. With probability at least $1-e^{-\tau}$,*

$$||\hat{\mu} - \mu^\star||_q \leq \min\left\{ 8\sqrt{\frac{\tau+m}{N}}, \ \frac{m^{1/q}}{\sqrt{2}}\sqrt{\frac{\tau+\log(2m)}{N}} \right\}.$$

The second term in the above min is immediate from applying Hoeffding's inequality to each component and using the union bound. For $q=1$ and any τ, however, the Hoeffding-based bound diverges for any $N=O(m^2)$ (in general it goes to zero only for $N=\omega(m^{2/q})$). We derive the first term using a Rademacher complexity argument. The first term goes to zero as long as $N=\omega(m)$ grows superlinearly in m. This first term is critical to show $\mu^\star$ can be estimated consistently in the q norm for $q<2$.

2.3. A Factored Gradient Descent Algorithm

In this section, we develop a specialized first-order algorithm for computing $\widehat{\Theta}$ that works on the non-convex, factored form of the nuclear-norm regularized likelihood optimization problem. The algorithm is particularly economical with memory because (a) it does not keep all $m \times n$ optimization variables in memory but rather only $\tilde{r} \times (m+n)$ where $\tilde{r}$ is a guess at rank and (b) it eschews the use of *any* spectral computation such as SVD (or partial SVD) at each iteration. This makes the algorithm particularly useful for the large scale data encountered in e-commerce applications.

Our factored gradient descent (FGD) algorithm solves the problem

$$\text{minimize } L(\Theta) + \lambda \|\Theta\|_*. \tag{6}$$

As discussed in Sec. 2.2, the algorithm we employ does *not* enforce any constraint on $||\Theta||_\infty$. A constraint of this form is necessary for the technical result in our main theorem, but is unnecessary in practice, as can be seen in our numerical results in Section 5.

In applications, one is interested in solving the problem (6) for very large m, n, N. Due to the complexity of Cholesky factorization, this rules out theoretically-tractable second-order interior point methods. One standard approach is to use a first-order method, such as Cai et al. (2010), Parikh and Boyd (2014), Hazan (2008), Orabona et al. (2012); however, this approach requires (at least a partial) SVD at each step. An alternative approach, which we take here, is to optimize as variables the factors $U \in \mathbb{R}^{m\times\tilde{r}}$ and $V \in \mathbb{R}^{n\times\tilde{r}}$ of the optimization variable $\Theta = UV^T$ rather than producing these via SVD at each step; see, *e.g.*, Keshavan et al. (2009b), Jain et al. (2013). To guarantee equivalence of the problems, we must take $\tilde{r} = \min(m,n)$. However, if we believe the solution is low rank, we may use a smaller $\tilde{r}$, reducing computational work and storage.

Our FGD algorithm proceeds by applying gradient descent steps to the unconstrained problem

$$\begin{array}{l}\text{minimize } L(UV^T) + \frac{\lambda}{2}\|U\|_{\mathrm{F}}^2 + \frac{\lambda}{2}\|V\|_{\mathrm{F}}^2, \\ \text{subject to } U \in \mathbb{R}^{m\times\tilde{r}},\ V \in \mathbb{R}^{n\times\tilde{r}}.\end{array} \tag{7}$$

This formulation has the advantage of being unconstrained, with a differentiable objective function, making it amenable to solution via simple optimization methods (Recht et al. 2010, 2011, Udell et al. 2016).

Algorithm 1 Factored Gradient Descent for (6)

input: dimensions m, n, $\tilde{r}$, data $\{(i_t, j_t, S_t)\}_{t=1}^N$, regularizing coefficient λ, and tolerance τ

$U \leftarrow U^0$, $V \leftarrow V^0$, $f' \leftarrow \infty$.

repeat

$\eta \leftarrow 1$, $f \leftarrow f'$, $\Delta U \leftarrow -\lambda U$, $\Delta V \leftarrow -\lambda V$

for $t = 1, \ldots, N$ **do**

$Q = 0$

for $j \in S_t$ **do**

$w_j \leftarrow e^{-U_{i_t}^T V_j}$

$W \leftarrow W + w_j$

end for

$\Delta U \leftarrow \Delta U - \frac{1}{N}\left(e_{i_t} V_{j_t}^T - \frac{1}{W}\sum_{j \in S_t} w_j e_{i_t} V_j^T\right)$

$\Delta V \leftarrow \Delta V - \frac{1}{N}\left(e_{j_t} U_{i_t}^T - \frac{1}{W}\sum_{j \in S_t} w_j e_j U_{i_t}^T\right)$

end for

repeat

$U' \leftarrow U + \eta \Delta U$, $V' \leftarrow V + \eta \Delta V$

$f' \leftarrow L(U'V'^T) + \frac{\lambda}{2}\|U'\|_{\mathrm{F}}^2 + \frac{\lambda}{2}\|V'\|_{\mathrm{F}}^2$

$\eta \leftarrow \beta_{\mathrm{dec}} \eta$

until $f' \leq f$

$U \leftarrow U'$, $V \leftarrow V'$

until $\frac{f - f'}{f'} \leq \tau$

output: UV^T

LEMMA 1. *Problem* (7) *is equivalent to*

$$\begin{array}{l} \text{minimize } L(\Theta) + \lambda \|\Theta\|_* \\ \text{subject to } \operatorname{rank}(\Theta) \leq \tilde{r}. \end{array} \tag{8}$$

That is, Problem (7) is equivalent to Problem (6) subject to an additional rank constraint $\operatorname{rank}(\Theta) \leq \tilde{r}$. If, for a particular choice of loss function L, Problem (6) has a solution with rank less than $\tilde{r}$, then the rank constraint is not binding, so Problem (6) is itself equivalent to Problem (7). (Recall that we defer all proofs to the Appendix.)

It is easy to compute the the gradients of the objective of (7). Since $L(\Theta)$ is differentiable,

$$\nabla_U L(UV^T) = \nabla L(UV^T) V,$$
$$\nabla_V L(UV^T) = \nabla L(UV^T)^T U.$$

We do not need to explicitly form $\nabla L(UV^T)$ in order to compute these; computing gradients implicitly reduces the memory required to implement the algorithm (see Algorithm 1). Similarly, we need not form UV^T to compute $L(UV^T)$. Recent work has shown that gradient descent on the factors converges linearly to the global optimum for problems that enjoy restricted strong convexity (Bhojanapalli et al. 2015). In eq. (12) in the proof of Theorem 3 we establish restricted strong convexity for our problem with high probability. Hence with high probability, FGD converges to the global minimum of Problem (7), and hence to the global minimum of Problem (6) provided $\tilde{r}$ is chosen to be large enough.

We initialize our algorithm using a technique from Bhojanapalli et al. (2015), which only requires access to gradients of the objective of (6). Using the SVD, we write $-\nabla L(0) = \tilde{U}\,\mathrm{diag}(\tilde{\sigma}_1,\dots,\tilde{\sigma}_{\min(m,n)})\tilde{V}^T$ and initialize

$$\begin{aligned} U^0 &= \gamma^{-1/2}\,\mathrm{diag}(\sqrt{\tilde{\sigma}_1},\dots,\sqrt{\tilde{\sigma}_{\tilde{r}}})\tilde{U}_{:,\tilde{r}}, \\ V^0 &= \gamma^{-1/2}\,\mathrm{diag}(\sqrt{\tilde{\sigma}_1},\dots,\sqrt{\tilde{\sigma}_{\tilde{r}}})\tilde{V}_{:,\tilde{r}}, \end{aligned}$$

where $\gamma = \|\nabla L(0) - (\nabla L(e_1 e_1^T) + \lambda e_1 e_1^T)\|_{\mathrm{F}}$ and $\tilde{U}_{:,\tilde{r}}$, $\tilde{V}_{:,\tilde{r}}$ denote the first r columns of $\tilde{U}$, $\tilde{V}$. We use an adaptive step size with a line search that guarantees descent. Starting with a stepsize of $\eta = 1$, the stepsize is repeatedly decreased by a factor β_{dec} until the step produces a decrease in the objective. We terminate the algorithm when the decrease in the relative objective value is smaller than the convergence tolerance τ.

3. The Dynamic Assortment Personalization Problem

We now return to the dynamic setting. First, we define some notation that will be useful in the following discussion. For $\theta \in \mathbb{R}^n$, $w \in \mathbb{R}^n$, $S \subset \{1,\dots,n\}$, and $K \le n$, let

$$\begin{aligned} p_j(S;\theta) &= \frac{e^{\theta_j}}{1+\sum_{j'\in S} e^{\theta_{j'}}}, && \text{the MNL choice probability under parameters } \theta, \\ F(S;w,\theta) &= \sum_{j\in S} p_j(S;\theta) w_j, && \text{the expected revenue of assortment } S \text{ under } \theta, \text{ and} \\ S^\star(w,\theta;K) &= \operatorname*{argmax}_{|S|\le K} F(S;w,\theta), && \text{the optimal assortment under } \theta \text{ for revenues } w. \end{aligned}$$

3.1. Structure-ignorant algorithms

An algorithm is *structure-ignorant* if it ignores any potential structure that connect the different contexts (rows of $\Theta^\star$). Therefore, a structure-ignorant algorithm is one that runs separate, independent algorithms for each context. Formally, we make the following definition.

Definition 2. An algorithm π for the dynamic assortment personalization problem is *structure-ignorant* if, under π, the event $S_t \mid i_t = i$ is measurable with respect to only the variables $\{i_{t'}, j_{t'}, S_{t'} : t' < t, i_{t'} = i\}$ and $U_1,\dots,U_t$.

A sequence of algorithms indexed by T is said to be consistent if it has sublinear regret in T over all problem instances.

DEFINITION 3. Fix (m,n,W). A sequence of algorithms π_T is said to be *consistent* if for any $\mu^\star$, $\Theta^\star$, and $a>0$, we have $\mathrm{Regret}(\pi_T)=o(T^a)$ over all problem instances $(m,n,T,W,\mu^\star,\Theta^\star)$.

For brevity, we usually omit the subscript and abuse notation in referring to a family or sequence of algorithms simply as one algorithm π.

If we run separate algorithms for each context then, in each context, we are solving the classic (non-contextual) dynamic assortment planning problem, precisely as studied by Rusmevichientong et al. (2010), Sauré and Zeevi (2013). Making use of Theorem 1 of Sauré and Zeevi (2013) we establish the following:

THEOREM 5. *Fix $\nu\in(0,1)$, $w\geq 0$, $\rho\geq 1$. Suppose that $T=\Omega(m^{1+\epsilon})$ grows superlinearly in the number of types m. Given any family (indexed by revealed problem parameters m,n,T,W) of structure-ignorant consistent algorithms π, we have*

$$\mathrm{Regret}(\pi)=\Omega(\min(m,n)\max(m,n)\log T)$$

over all the problem instances $(m,n,T,W,\mu^\star,\Theta^\star)$ satisfying, for all $i=1,\dots,m$:

1. *Profit is bounded:*

$$\|W_i\|_\infty\leq w.$$

2. *Each type appears often enough and not too often:*

$$1/\rho\leq m\mu_i\leq\rho.$$

3. *The number of* potentially optimal items *grows linearly in n:*

$$\left|\left\{j\notin S^\star(W_i,\Theta_i^\star;K):\ \exists\theta\in\mathbb{R}^n,\, j\in S^\star(W_i,\theta;K),\,\theta_{j'}=\Theta_{ij}^\star\ \forall j'\in S^\star(W_i,\Theta_i^\star;K)\right\}\right|\geq\nu n.$$

Note that $mn=\min(m,n)\max(m,n)$. We rewrite the product mn in this way to make the comparison with our structure-aware algorithms more clear: they depend on $\max(m,n)$, but replace the $\min(m,n)$ term by the rank r of the parameter matrix Θ.

Theorem 5 asserts that structure-ignorant algorithms for the dynamic assortment personalization problem have regret that grows linearly in the number of item-type combinations and logarithmically in the time horizon. That is, ignoring structure incurs *enormous* regret. In the sequel, we improve on this regret bound by developing a structure-aware algorithm.

Algorithm 2 Dynamic Assortment Personalization via Nuclear-Norm Regularization ($\pi_{\text{nuc-norm}}(T_0, \lambda)$)

input: T_0, λ

Exploration stage:

for $t = 1, \dots, T_0$ **do**

Let S_t be drawn uniformly at random from all subsets of size K.

end for

Estimation:

Define

$$L(\Theta) = \frac{1}{T_0} \sum_{t=1}^{T_0} \log \left(\left(1 + \sum_{j \in S_t} e^{\Theta_{ij}} \right) \left(\begin{cases} 1 & j_t = 0 \\ e^{\Theta_{ij_t}} & \text{otherwise} \end{cases} \right)^{-1} \right).$$

Solve

$$\widehat{\Theta} \in \operatorname{argmin} \; L(\Theta) + \lambda \|\Theta\|_* \quad \text{s.t. } \|\Theta\|_\infty \leq \alpha / \sqrt{mn}.$$

Exploitation stage:

for $t = T_0 + 1, \dots$ **do**

Let

$$S_t \in S^\star(W_{i_t}, \widehat{\Theta}_{i_t}; K).$$

end for

output: $S_1, S_2, \dots$

3.2. A Structure-Aware Algorithm

In Section 2, we argued that imposing structure is crucial for learning the preferences of a very heterogeneous population and proposed an estimator that can leverage structure to learn the LRMMNL model in sublinear time. In this section, we use this estimator to develop a structure-aware algorithm that can achieve regret sublinear in problem size mn.

In Algorithm 2, we define an algorithm for dynamic assortment personalization with tuning parameters T_0 and λ. We refer to this algorithm as $\pi_{\text{nuc-norm}}$.

Every step of Algorithm 2 is computationally tractable, including the computation of the optimal assortment $S^\star(W_{i_t}, \widehat{\Theta}_{i_t}; K)$ (Rusmevichientong et al. 2010) and of $\widehat{\Theta}$ (Algorithm 1). Since $S^\star(W_i, \widehat{\Theta}_i; K)$ can be computed once, for each i, at the start of the exploitation stage, this algorithm can be run in time polynomial in m, n, and T_0.

Next we show that this achieves regret that is order-of-magnitude smaller than a structure-ignorant algorithm when structure is present. Let

$$\delta(w,\theta;K) = \max_{|S|\le K} F(S(w,\theta;K);w,\theta) - \max_{S\notin S^\star(w,\theta;K):|S|\le K} F(S;w,\theta)$$

be the gap in revenue between the optimal assortment and any suboptimal assortment under MNL choice with parameter vector θ.

THEOREM 6. *Let $\rho \ge 1$ be such that $1/\rho \le m\mu_i \le \rho\ \forall i = 1,\dots,m$, α be such that $||\Theta^\star||_\infty \le \alpha/\sqrt{mn}$, δ be such that $\min_{1\le i\le m}\delta(W_i,\Theta_i^\star;K)\ge\delta$, and ω be such that $||W||_\infty \le \omega$. Let $C = 1048576K^6\rho^3\omega^2\alpha^2\exp(8\alpha)/\delta^2$ be a constant defined in terms of the given constants and let*

$$T_0 = Cr(m+n)\log(T), \quad \lambda = 8\sqrt{\frac{\rho K}{Crmn}}.$$

Suppose $T \le (m+n)^{\frac{mn}{c(m+n}r}$. Then the regret of $\pi_{nuc\text{-}norm}(T_0,\lambda)$ satisfies

$$\begin{aligned}\text{Regret}\left(\pi_{nuc\text{-}norm}(T_0,\lambda)\right) &\le 3\omega + C\omega r(m+n)\log(T)\\ &= O\left(r\max(m,n)\log(T)\right).\end{aligned}$$

This shows that leveraging the structure appropriately can lead to regret that is potentially an order of magnitude smaller. That is, if rank r is constant, or grows sublinearly (*e.g.*, logarithmically) in the side dimension $\min(m,n)$, then the regret of the $\pi_{\text{nuc-norm}}$ algorithm also grows sublinearly. That is,

$$\text{Regret}\left(\pi_{\text{nuc-norm}}(T_0,\lambda)\right) = O(r\max(m,n)\log(T)).$$

This regret is (asymptotically) smaller than $\min(m,n)\max(m,n)\log(T)$, the lower bound on the regret of any structure-ignorant consistent algorithm, by a factor of $\min(m,n)/r$.

4. The Dynamic Assortment Planning Problem with a Heterogeneous Population

In this section we consider a non-contextual dynamic assortment planning problem where the target population is very heterogeneous. This is appropriate for settings without the potential for personalization, such as a single brick-and-mortar store, but where the consumer population is very heterogeneous and cannot be well described by only a single, or even a few, nominal preference vectors. This setting is similar to that considered in Rusmevichientong et al. (2010) and Sauré and Zeevi (2013) in that it is non-contextual, but in both those works the choice distribution is assumed to be MNL (or, generally, any random utility model with a single nominal preference vector, from which deviations are made at random in a homogeneous manner). Here we allow for a potentially heterogeneous population with choice governed by the LRMMNL model.

The problem proceeds as follows, in somewhat different order than the personalization problem presented in Section 1.1. At each $t = 1,\dots,T$:

1. we choose any $S_t \subset \{1,\dots,n\}$ with $|S_t| \le K$,
2. a type i_t is drawn at random from $\{1,\dots,m\}$ with probability proportional to weights $\mu_i^\star$,
3. an item j_t is drawn at random from $\{1,\dots,n\}$ with probability proportional to weights

$$\text{weight}(j) = \begin{cases} 1 & j = 0 \\ 0 & j \neq 0,\ j \notin S_t \\ \exp(\Theta^\star_{i_t j}) & j \neq 0,\ j \in S_t \end{cases}$$

4. if $j_t = 0$ we get reward $r_t = 0$ and otherwise we get reward $r_t = W_{i_t j_t}$.

Unlike before, we have to select S_t *before* observing i_t. Therefore, we cannot personalize.

Let us re-define

$$F(S; W, \Theta, \mu) = \sum_{i=1}^m \mu_i \frac{\sum_{j\in S} e^{\Theta_{ij}} W_{ij}}{1 + \sum_{j\in S} e^{\Theta_{ij}}}, \qquad \text{the expected revenue of } S \text{ under } \Theta, \text{ and}$$

$$S^\star(W, \Theta, \mu; K) = \arg\max_{|S|\le K} F(S; W, \Theta, \mu), \qquad \text{the optimal assortment under } \Theta.$$

Note that $S^\star(W, \Theta, \mu; K)$ is not efficiently computable; however, an efficient approximation scheme, which searches over revenue-ordered assortments, is proposed in Rusmevichientong et al. (2014) and shown to work well.

We adapt our nuclear-norm-regularized algorithm to this case as shown in Algorithm 3. We refer to this algorithm as $\pi_{\text{nuc-norm-plan}}$.

Next we show that this achieves regret that is order-of-magnitude smaller than a structure-ignorant algorithm when structure is present. Let

$$\delta(W, \Theta, \mu; K) = \max_{|S|\le K} F(S; W, \Theta, \mu) - \max_{S \notin S^\star(W,\Theta,\mu;K): |S|\le K} F(S; W, \Theta, \mu)$$

be the gap in revenue between the optimal assortment and any suboptimal assortment under parameter matrix Θ.

THEOREM 7. *Let $\rho \ge 1$ be such that $1/\rho \le m\mu_i \le \rho\ \forall i = 1,\dots,m$, α be such that $||\Theta^\star||_\infty \le \alpha/\sqrt{mn}$, δ be such that $\min_{1\le i\le m} \delta(W_i, \Theta_i^\star; K) \ge \delta$, and ω be such that $||W||_\infty \le \omega$. Let $C = 1048576K^6\rho^3\omega^2\alpha^2\exp(8\alpha)/\delta^2$ be a constant defined in terms of the given constants and let*

$$T_0 = Cr(1 + \frac{n}{m})\log(T), \quad \lambda = 8\sqrt{\frac{\rho K}{Crmn}}.$$

Then the regret of $\pi_{nuc\text{-}norm\text{-}plan}(T_0, \lambda)$ satisfies

$$\begin{aligned} \text{Regret}\left(\pi_{nuc\text{-}norm\text{-}plan}(T_0, \lambda)\right) &\le 4\omega + C\omega r(m+n)\log(T) \\ &= O\left(r\max(m,n)\log(T)\right). \end{aligned}$$

5. Experimental Results

In this section we demonstrate numerically the importance of structure and the power of our approach.

Algorithm 3 Dynamic Assortment Planning via Nuclear-Norm Regularization ($\pi_{\text{nuc-norm-plan}}(T_0, \lambda)$)

input: T_0, λ

Exploration stage:

for $t = 1, \dots, T_0$ **do**

Let S_t be drawn uniformly at random from all subsets of size K.

end for

Estimation:

Define

$$L(\Theta) = \frac{1}{T_0} \sum_{t=1}^{T_0} \log \left(\left(1 + \sum_{j \in S_t} e^{\Theta_{ij}} \right) \left(\begin{cases} 1 & j_t = 0 \\ e^{\Theta_{ij_t}} & \text{otherwise} \end{cases} \right)^{-1} \right).$$

Solve

$$\widehat{\Theta} \in \operatorname{argmin} L(\Theta) + \lambda \|\Theta\|_* \quad \text{s.t. } \|\Theta\|_\infty \le \alpha / \sqrt{mn}$$

Set

$$\hat{\mu}_i = \frac{1}{T_0} \sum_{t=1}^{T_0} \mathbb{I}[i_t = i] \text{ for } i = 1, \dots, m$$

Exploitation stage:

for $t = T_0 + 1, \dots$ **do**

Choose action

$$S_t \in S^\star(W, \widehat{\Theta}, \hat{\mu}; K).$$

end for

output: $S_1, S_2, \dots$

Figure 1 **RMSE of estimators for $\Theta^\star$ with $m = n = 100, 150, 200, \dots, 500, 750, \dots, 4000$ and $r = 2, 25$.**

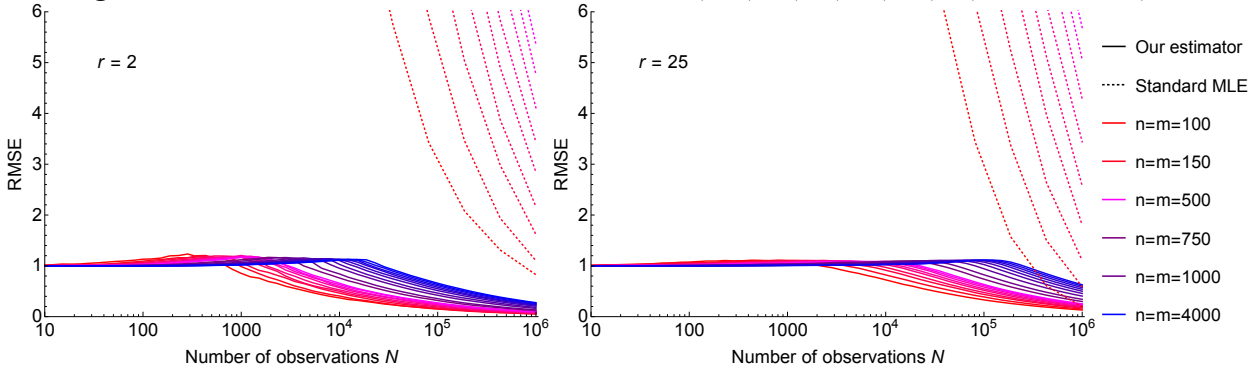

5.1. The Static Estimation Problem

First, we focus on the static estimation problem. We compare our estimate $\widehat{\Theta}$ with the standard maximum likelihood estimate $\widehat{\Theta}^{\mathrm{MLE}}$ that solves

$$\text{minimize } L(\Theta). \tag{9}$$

Note that, since it imposes no structure on the whole matrix Θ, problem (9) decomposes into m subproblems for each type (row of $\widehat{\Theta}^{\mathrm{MLE}}$), each solving a separate MNL MLE in n variables. In our experiments, we use Newton's method as implemented by `Optim.jl` to solve each subproblem.

To generate $\Theta^\star$, we fix m, n, r, let Θ_0 be an $m \times n$ matrix composed of independent draws from a standard normal, take its SVD $\Theta_0 = U \operatorname{diag}(\sigma_1, \sigma_2, \dots) V^T$, truncate it past the top r components $\Theta_1 = U \operatorname{diag}(\sigma_1, \dots, \sigma_r, 0, \dots) V^T$, and renormalize to achieve unit sample standard deviation to get $\Theta^\star$, *i.e.*, $\Theta^\star = \Theta_1 / \operatorname{std}(\operatorname{vec}(\Theta_1))$. To generate the choice data, we let i_t be drawn uniformly at random from $\{1, \dots, m\}$, S_t be drawn uniformly at random from all subsets of size 10, and j_t be chosen according to (1) with parameter $\Theta^\star$.

For our estimator we use Algorithm 1 with $\tilde{r} = 2r$, $\lambda = \frac{1}{8}\sqrt{\frac{Kd\log d}{mnT}}$, $\beta_{\mathrm{dec}} = 0.8$, and $\tau = 10^{-10}$. This regularizing coefficient scales with m, n, d, T, and K as suggested by Theorem 3, but we find the algorithm performs better in practice when we use a smaller constant than that suggested by the theorem.

We plot the results in Figure 1, where error is measured in root mean squared error (RMSE)

$$\mathrm{RMSE}(\Theta) = \sqrt{\mathrm{Avg}\left(\{(\Theta_{ij} - \Theta^\star_{ij})^2\}_{i,j}\right)} = \frac{1}{\sqrt{mn}} \left|\left|\Theta - \Theta^\star\right|\right|_{\mathrm{F}}.$$

The results show the advantage in efficient use of the data offered by our approach. The results also show that, relative to MLE, the advantage is greatest when the underlying rank r is small and the number of parameters $m \times n$ is large, but that we maintain a significant advantage even for moderate r and $m \times n$. For large numbers of parameters ($m = n \geq 750$), the RMSE of MLE is very large and does not appear in the plots. Only in the case of greatest rank ($r = 25$), smallest number of parameters ($m = n = 100$), and greatest number of observations ($T = 10^6$) does MLE appear to somewhat catch up with our estimator.

In Figure 2, we plot the RMSE of our estimator against the number of observations per type (or item) T/d for a square problem with $d = m = n$. We see nearly the same error curve traced out as we vary the problem size d. This scaling shows that our estimator is able to leverage the low-rank assumption and require the same number of choice observations per type to achieve the same RMSE regardless of problem size.

Figure 2 **RMSE of our estimators for $\Theta^\star$ by observation per row with $r = 2$.**

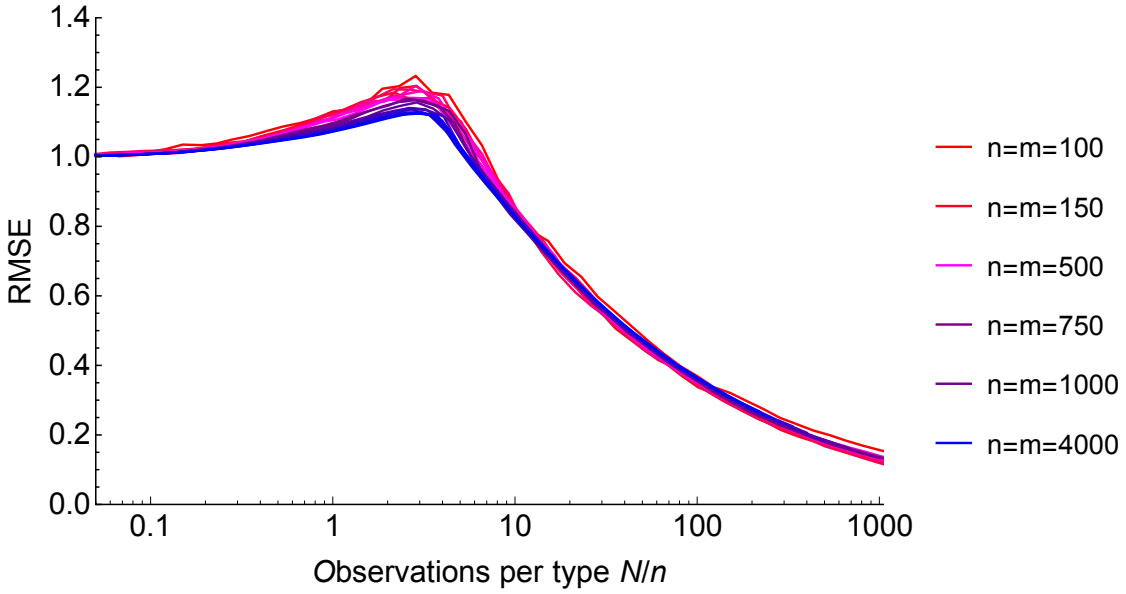

5.2. The Dynamic Assortment Personalization Problem

Next, we turn our attention to the dynamic assortment personalization problem. We compare our algorithm $\pi_{\text{nuc-norm}}$ to two alternatives. One alternative is the structure-ignorant algorithm in which we apply Algorithm 1 of Sauré and Zeevi (2013) to each type separately. Recognizing that one cannot learn a huge set of parameters from very few observations, the second alternative, which we term context-ignorant, tries to fit a single MNL model to the whole population, applying only one replicate Algorithm 1 of Sauré and Zeevi (2013) to all types simultaneously. We generate the problem exactly as in the previous section and run each of these algorithms on problem dimensions varying from 10,000 to 160,000 and for $r = 2$. We plot the regret over time in Figure 3.

Since the plots have a logarithmic horizontal axis, regret that is logarithmic in T appears as a line in the figure whereas regret that is linear in T appears as an exponential function. These plots reveal several interesting features of the algorithms.

We see that the structure-ignorant algorithm, which we know has logarithmic regret asymptotically, only achieves logarithmic regret for the smallest of the problems ($m = n = 100$). For larger problem sizes, regret appears to be linear for all horizons T shown. The transition from linear to logarithmic regret is not visible because the problems are so large that the transition occurs for extremely large T: on the scale of this plot, the mn term overwhelms the $\log(T)$ term.

In all of these larger problems, the context-ignorant algorithm performs better than the structure-ignorant algorithm. Both have linear regret in this parameter regime. In fact, the context-ignorant algorithm uses a misspecified model, and so has asymptotically linear regret. Hence for very large T (not shown on this plot) it will be overtaken by the structure-ignorant algorithm, whose regret

Figure 3 **Regret for Structure-Aware, Structure-Ignorant, and Context-Ignorant Algorithms (note horizontal log scale)**

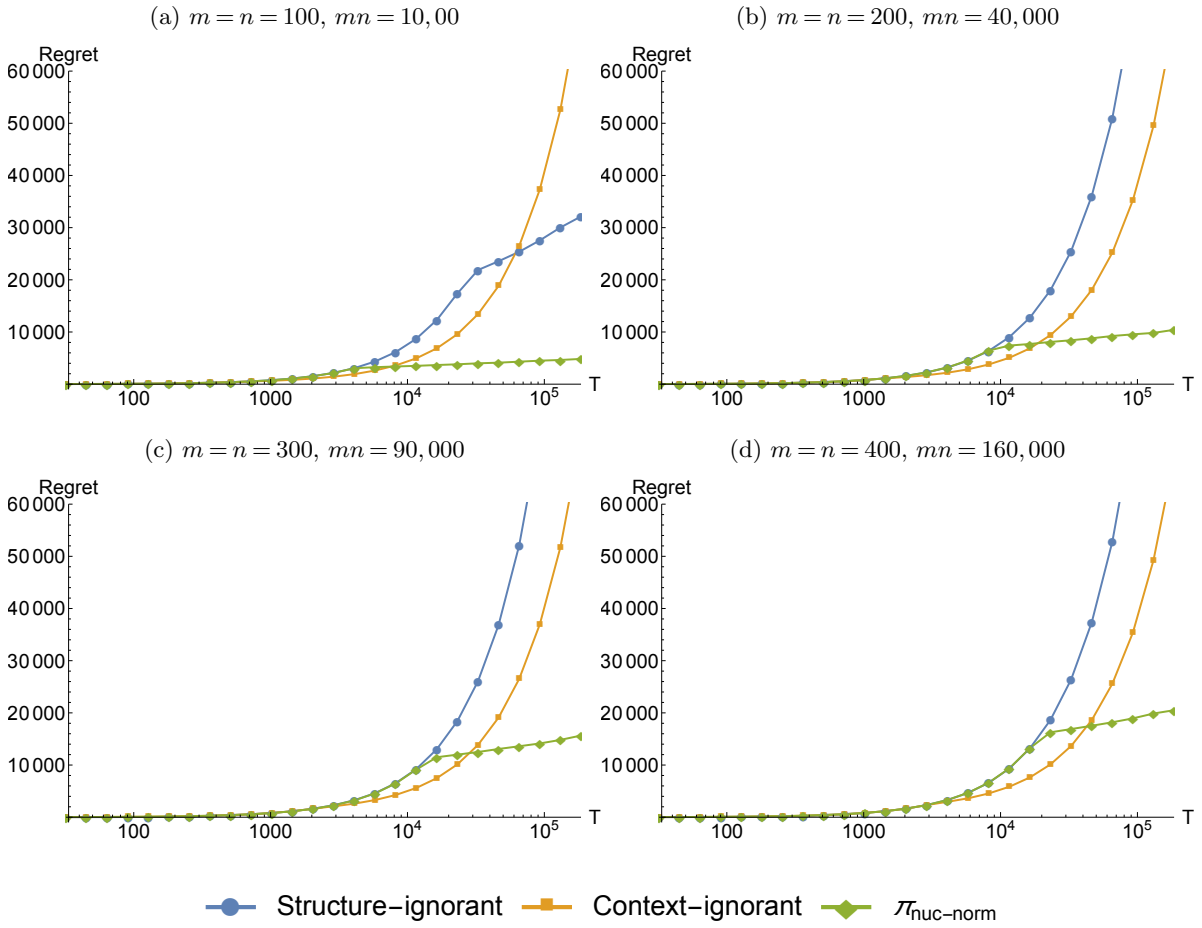

is asymptotically logarithmic. The success of the misspecified context-ignorant algorithm holds an important lesson: when time is limited, it is more effective to use a misspecified model with few parameters than a well specified model with many parameters.

On the other hand, our structure-aware algorithm exhibits logarithmic regret in each and every case, even at this relatively short time scale. The algorithm $\pi_{\text{nuc-norm}}$, as promised by Theorem 6, has logarithmic regret that does not explode astronomically with m, n.

6. Conclusion

To manage revenue, many retailers must solve a dynamic assortment personalization problem: they must learn customers preferences in real time, at scale, from customers' choices from among the items on offer; and they must quickly use this information to present revenue-maximizing assortments. This paper explores a structural approach to enable large scale dynamic assortment

personalization. We proposed algorithms using structural (low rank) priors to learn and exploit customer preferences. We presented theoretical and numerical evidence that these algorithms improve on the state of the art by orders of magnitude, and achieve performance suitable for use in practice.

References

Bartlett, Peter L, Shahar Mendelson. 2003. Rademacher and gaussian complexities: Risk bounds and structural results. *The Journal of Machine Learning Research* **3** 463–482.

Bastani, Hamsa, Mohsen Bayati. 2015. Online decision-making with high-dimensional covariates. *Available at SSRN 2661896* .

Bernstein, Fernando, A Gürhan Kök, Lei Xie. 2011. Dynamic assortment customization with limited inventories. Tech. rep., Citeseer.

Bertsimas, Dimitris, Nathan Kallus. 2014. From predictive to prescriptive analytics. *arXiv preprint arXiv:1402.5481* .

Bhojanapalli, Srinadh, Anastasios Kyrillidis, Sujay Sanghavi. 2015. Dropping convexity for faster semi-definite optimization. *arXiv preprint arXiv:1509.03917* .

Blanchet, Jose, Guillermo Gallego, Vineet Goyal. 2013. A Markov chain approximation to choice modeling. *EC*. 103–104.

Bubeck, Sébastien, Nicolo Cesa-Bianchi. 2012. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Machine Learning* **5**(1) 1–122.

Cai, Jian-Feng, Emmanuel J Candès, Zuowei Shen. 2010. A singular value thresholding algorithm for matrix completion. *SIAM Journal on Optimization* **20**(4) 1956–1982.

Candès, E, Y. Plan. 2009. Matrix completion with noise. *CoRR* **abs/0903.3131**.

Candès, E., B. Recht. 2009. Exact matrix completion via convex optimization. *Foundations of Computational Mathematics* **9**(6) 717–772.

Candès, E., T. Tao. 2010. The power of convex relaxation: Near-optimal matrix completion. *IEEE Transactions on Information Theory* **56**(5) 2053–2080.

Caro, Felipe, Jérémie Gallien. 2007. Dynamic assortment with demand learning for seasonal consumer goods. *Management Science* **53**(2) 276–292.

Davenport, M., Y. Plan, E. van den Berg, M. Wootters. 2014. 1-bit matrix completion. *Information and Inference* **3**(3) 189–223.

Davis, James M, Guillermo Gallego, Huseyin Topaloglu. 2014. Assortment optimization under variants of the nested logit model. *Operations Research* **62**(2) 250–273.

DeShazo, JR, German Fermo. 2002. Designing choice sets for stated preference methods: the effects of complexity on choice consistency. *Journal of Environmental Economics and management* **44**(1) 123–143.

Désir, Antoine, Vineet Goyal. 2014. Near-optimal algorithms for capacity constrained assortment optimization. *Available at SSRN 2543309* .

Farias, Vivek F, Srikanth Jagabathula, Devavrat Shah. 2013. A nonparametric approach to modeling choice with limited data. *Management Science* **59**(2) 305–322.

Goldenshluger, Alexander, Assaf Zeevi. 2013. A linear response bandit problem. *Stochastic Systems* **3**(1) 230–261.

Golrezaei, Negin, Hamid Nazerzadeh, Paat Rusmevichientong. 2014. Real-time optimization of personalized assortments. *Management Science* **60**(6) 1532–1551.

Gunasekar, S., P. Ravikumar, J. Ghosh. 2014. Exponential family matrix completion under structural constraints. *Proceedings of the 31st International Conference on Machine Learning (ICML-14)*. 1917–1925.

Hazan, Elad. 2008. Sparse approximate solutions to semidefinite programs. *LATIN 2008: Theoretical Informatics*. Springer, 306–316.

Jain, Prateek, Praneeth Netrapalli, Sujay Sanghavi. 2013. Low-rank matrix completion using alternating minimization. *Proceedings of the forty-fifth annual ACM Symposium on the Theory of Computing*. ACM, 665–674.

Kahneman, Daniel, Amos Tversky. 1979. Prospect theory: An analysis of decision under risk. *Econometrica: Journal of the Econometric Society* 263–291.

Kallus, Nathan, Madeleine Udell. 2016. Revealed preference at scale: Learning personalized preferences from assortment choices. *Proceedings of the 2016 ACM Conference on Economics and Computation*. ACM, 821–837.

Keshavan, R., A. Montanari, S. Oh. 2009a. Matrix completion from noisy entries. *Advances in Neural Information Processing Systems*. 952–960.

Keshavan, R. H., A. Montanari, S. Oh. 2010. Matrix completion from a few entries. *IEEE Transactions on Information Theory* **56**(6) 2980–2998.

Keshavan, Raghunandan, Andrea Montanari, Sewoong Oh. 2009b. Matrix completion from noisy entries. *Advances in Neural Information Processing Systems*. 952–960.

Lai, Tze Leung, Herbert Robbins. 1985. Asymptotically efficient adaptive allocation rules. *Advances in applied mathematics* **6**(1) 4–22.

Li, Guang, Paat Rusmevichientong, Huseyin Topaloglu. 2015. The d-level nested logit model: Assortment and price optimization problems. *Operations Research* **63**(2) 325–342.

Lu, Yu, Sahand N Negahban. 2014. Individualized rank aggregation using nuclear norm regularization. *arXiv preprint arXiv:1410.0860* .

Luce, Duncan R. 1959. *Individual Choice Behavior: a theoretical analysis*. Wiley.

McFadden, Daniel. 1973. Conditional logit analysis of qualitative choice behavior. P. Zarembka, ed., *Frontiers in Economics*. Academic Press.

McFadden, Daniel. 1980. Econometric models for probabilistic choice among products. *Journal of Business* S13–S29.

McFadden, Daniel, Kenneth Train. 2000. Mixed mnl models for discrete response. *Journal of applied Econometrics* **15**(5) 447–470.

Megiddo, Nimrod. 1979. Combinatorial optimization with rational objective functions. *Mathematics of Operations Research* **4**(4) 414–424.

Negahban, Sahand, Martin J Wainwright. 2011. Estimation of (near) low-rank matrices with noise and high-dimensional scaling. *The Annals of Statistics* 1069–1097.

Oh, Sewoong, Kiran K Thekumparampil, Jiaming Xu. 2015. Collaboratively learning preferences from ordinal data. *arXiv preprint arXiv:1506.07947* .

Orabona, Francesco, Andreas Argyriou, Nathan Srebro. 2012. Prisma: Proximal iterative smoothing algorithm. *arXiv preprint arXiv:1206.2372* .

Parikh, Neal, Stephen Boyd. 2014. Proximal algorithms. *Foundations and Trends in Optimization* **1**(3) 127–239.

Perchet, Vianney, Philippe Rigollet. 2013. The multi-armed bandit problem with covariates. *The Annals of Statistics* **41**(2) 693–721.

Recht, B., M. Fazel, P. Parrilo. 2010. Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization. *SIAM Review* **52**(3) 471–501. doi:10.1137/070697835. URL `http://dx.doi.org/10.1137/070697835`.

Recht, Benjamin, Christopher Re, Stephen Wright, Feng Niu. 2011. Hogwild: A lock-free approach to parallelizing stochastic gradient descent. *Advances in Neural Information Processing Systems*. 693–701.

Rigollet, Philippe, Assaf Zeevi. 2010. Nonparametric bandits with covariates. *arXiv preprint arXiv:1003.1630* .

Rusmevichientong, Paat, Zuo-Jun Max Shen, David B Shmoys. 2010. Dynamic assortment optimization with a multinomial logit choice model and capacity constraint. *Operations Research* **58**(6) 1666–1680.

Rusmevichientong, Paat, David Shmoys, Chaoxu Tong, Huseyin Topaloglu. 2014. Assortment optimization under the multinomial logit model with random choice parameters. *Production and Operations Management* **23**(11) 2023–2039.

Sauré, Denis, Assaf Zeevi. 2013. Optimal dynamic assortment planning with demand learning. *Manufacturing & Service Operations Management* **15**(3) 387–404.

Slivkins, Aleksandrs. 2014. Contextual bandits with similarity information. *The Journal of Machine Learning Research* **15**(1) 2533–2568.

Talluri, Kalyan T, Garrett J Van Ryzin. 2006. *The theory and practice of revenue management*, vol. 68. Springer Science & Business Media.

Tropp, Joel A. 2012. User-friendly tail bounds for sums of random matrices. *Foundations of Computational Mathematics* **12**(4) 389–434.

Udell, M., C. Horn, R. Zadeh, S. Boyd. 2016. Generalized low rank models. *Foundations and Trends in Machine Learning* URL `http://dx.doi.org/10.1561/2200000055`.

van Ryzin, Garrett, Gustavo Vulcano. 2014. A market discovery algorithm to estimate a general class of nonparametric choice models. *Management Science* **61**(2) 281–300.

Williams, Huw CWL. 1977. On the formation of travel demand models and economic evaluation measures of user benefit. *Environment and planning A* **9**(3) 285–344.

Appendices

7. Proof of offline recovery guarantee

First we introduce some notation and preliminary results. Define $\Delta = \widehat{\Theta} - \Theta^\star$ as the error to bound, $S_t' = S_t \backslash \{j_t\}$, $\gamma = \alpha/\sqrt{mn}$, $X_{tj} = e_{i_t} e_j^T$ where e_ℓ for $\ell \geq 1$ is the ℓ^{th} unit vector and e_0 is the vector of all zeros in the appropriate spaces ($e_{i_t} \in \mathbb{R}^m$, $e_{j_t} \in \mathbb{R}^n$), and $Y_t(\Theta) = \frac{1}{K_t} \sum_{j_t \in S_t} \Theta_{i_t j}^2$.

This notation implies

$$
\begin{aligned}
L(\Theta) &= \frac{1}{N} \sum_{t=1}^N \left(\log\left(1 + \sum_{j \in S_t} e^{X_{tj} \cdot \Theta}\right) - X_{tj_t} \cdot \Theta \right) \\
\nabla L(\Theta) &= \frac{1}{N} \sum_{t=1}^N \left(\frac{\sum_{j \in S_t} e^{X_{tj} \cdot \Theta} X_{tj}}{1 + \sum_{j \in S_t} e^{X_{tj} \cdot \Theta}} - X_{tj_t} \right) \\
\nabla^2 L(\Theta) &= \frac{1}{N} \sum_{t=1}^N \frac{\left(1 + \sum_{j \in S_t} e^{X_{tj} \cdot \Theta}\right) \left(\sum_{j \in S_t} e^{X_{tj} \cdot \Theta} X_{tj}^{\otimes 2}\right) - \left(\sum_{j \in S_t} e^{X_{tj} \cdot \Theta} X_{tj}\right)^{\otimes 2}}{\left(1 + \sum_{j \in S_t} e^{X_{tj} \cdot \Theta}\right)^2}
\end{aligned}
$$

where $A^{\otimes 2} = A \otimes A$ is the outer product power that yields a the symmetric linear operator on matrices $(A \otimes A)(B) = (A \cdot B)A$.

We will use the following three lemmas, whose proofs are presented after that of Theorem 3.

LEMMA 2. *Let* $\mathcal{A}_{\Gamma,\nu} = \{\Theta : \|\Theta\|_\infty \leq \gamma, \|\Theta\|_F \leq \Gamma, \|\Theta\|_* \leq \frac{\nu}{60\sqrt{\rho^3 K m n \gamma}} \sqrt{\frac{N}{(m+n)\log(m+n)}} \Gamma^2\}$ *and*

$$
\mathcal{M}_{\Gamma,\nu} = \sup_{\Theta \in \mathcal{A}_{\Gamma,\nu}} \left(\frac{1}{\rho m n} \|\Theta\|_F^2 - \frac{1}{N} \sum_{t=1}^N Y_t(\Theta) \right).
$$

Then

$$
\mathbb{P}\left(\mathcal{M}_{\Gamma,\nu} \geq \nu \frac{\Gamma^2}{\rho m n} \right) \leq \exp\left(-\frac{8}{9} \frac{\nu^2}{\rho^2 m^2 n^2} \frac{\Gamma^4}{\gamma^4} N \right)
$$

LEMMA 3. *Let* $\tau \geq 1$ *be given. Letting*

$$
\mathcal{A}^\star = \left\{ \Theta : \|\Theta\|_\infty \leq \gamma, \|\Theta\|_* \leq \frac{1}{\max\{(18\tau)^{1/4}, 480\}} \frac{1}{\sqrt{\rho^3 K m n}\gamma} \sqrt{\frac{N}{\sqrt{mn}\log(m+n)}} \|\Theta\|_F^2 \right\},
$$

we have

$$
\mathbb{P}\left(\frac{1}{N} \sum_{t=1}^N Y_t(\Theta) \geq \frac{1}{2\rho m n} ||\Theta||_F^2 \ \ \forall \Theta \in \mathcal{A}^\star \right) \geq 1 - 2(m+n)^{-\tau}.
$$

LEMMA 4. *Let* $\tau \geq 1$ *be given. With probability at least* $1 - (m+n)^{-\tau}$,

$$
||\nabla L(\Theta^\star)||_2 \leq 4\sqrt{\tau} \sqrt{\frac{\rho K (m+n)\log(m+n)}{mnN}}.
$$

Proof of Theorem 3 With probability at least $1-3(m+n)^{-\tau}$, the events in both Lemma 3 and Lemma 4 occur. Let us first assume $\Delta \in \mathcal{A}^\star$.

Define $D = L(\widehat{\Theta}) - L(\Theta^\star) - \nabla L(\Theta^\star)\cdot(\widehat{\Theta}-\Theta^\star)$. By Taylor's theorem, $\exists s\in[0,1]$ such that

$$\begin{aligned}
D &= \nabla^2 L(\Theta^\star + s\Delta)[\Delta,\Delta] \\
&= \frac{1}{N}\sum_{t=1}^N \frac{\left(1+\sum_{j\in S_t} e^{v_{tj}}\right)\left(\sum_{j\in S_t} e^{v_{tj}}(X_{tj}\cdot\Delta)^2\right) - \left(\sum_{j\in S_t} e^{v_{tj}} X_{tj}\cdot\Delta\right)^2}{\left(1+\sum_{j\in S_t} e^{v_{tj}}\right)^2} \\
&\geq \frac{1}{N}\sum_{t=1}^N \frac{\sum_{j\in S_t} e^{v_{tj}}(X_{tj}\cdot\Delta)^2}{\left(1+\sum_{j\in S_t} e^{v_{tj}}\right)^2} \geq \frac{1}{N}\frac{1}{e^{4\gamma}}\sum_{t=1}^N \frac{1}{(K_t+1)^2}\sum_{j\in S_t}(X_{tj}\cdot\Delta)^2 \\
&\geq \frac{1}{N}\frac{1}{e^{4\gamma}}\frac{1}{4K}\sum_{t=1}^N Y_t(\Delta), \qquad (10)
\end{aligned}$$

where $v_{tj} = X_{tj}\cdot(\Theta^\star + s\Delta)$, the first inequality is Jensen's, and the second is from $|v_{tj}|\leq 2\gamma$.

We now decompose Δ as in Recht et al. (2010), Negahban and Wainwright (2011). Let $\Theta^\star = U\,\mathrm{Diag}(\sigma_1,\sigma_2,\dots)V^T$ be the singular-value decomposition of $\Theta^\star$ with singular values sorted largest to smallest. Using block notation, let us write

$$U^T\Delta V = \Gamma = \begin{pmatrix}\Gamma_{11} & \Gamma_{12}\\ \Gamma_{21} & \Gamma_{22}\end{pmatrix} \text{ with } \Gamma_{11}\in\mathbb{R}^{r\times r}$$
$$\Delta'' = U\begin{pmatrix}0 & 0\\ 0 & \Gamma_{22}\end{pmatrix}V^T, \quad \Delta' = U\begin{pmatrix}\Gamma_{11} & \Gamma_{12}\\ \Gamma_{21} & 0\end{pmatrix}V^T.$$

Then $\Delta = U\Gamma V^T = \Delta' + \Delta''$. Note that,

$$\begin{aligned}
\mathrm{rank}(\Delta') &= \mathrm{rank}(U^T\Delta' V) \\
&= \mathrm{rank}\left(\begin{pmatrix}\Gamma_{11}/2 & \Gamma_{12}\\ 0 & 0\end{pmatrix} + \begin{pmatrix}\Gamma_{11}/2 & 0\\ \Gamma_{21} & 0\end{pmatrix}\right) \leq 2r.
\end{aligned}$$

Letting $\Theta^\star_r = U\,\mathrm{Diag}(\sigma_1,\dots,\sigma_r,0,0,\dots)V^T$ and its complement $\overline{\Theta}^\star_r = \Theta^\star - \Theta^\star_r$, we see

$$\begin{aligned}
\|\widehat{\Theta}\|_* &= \|\Theta^\star + \Delta\|_* = \|\Theta^\star_r + \overline{\Theta}^\star_r + \Delta' + \Delta''\|_* \\
&\geq \|\Theta^\star_r + \Delta''\|_* - \|\overline{\Theta}^\star_r\|_* - \|\Delta'\|_* \\
&= \|\Theta^\star_r\|_* + \|\Delta''\|_* - \|\overline{\Theta}^\star_r\|_* - \|\Delta'\|_* \\
&= \|\Theta^\star\|_* + \|\Delta''\|_* - 2\|\overline{\Theta}^\star_r\|_* - \|\Delta'\|_*,
\end{aligned}$$

and so $\|\widehat{\Theta}\|_* - \|\Theta^\star\|_* \leq 2\|\overline{\Theta}^\star_r\|_* + \|\Delta'\|_* - \|\Delta''\|_*$.

By the optimality of $\widehat{\Theta}$, we have

$$L(\widehat{\Theta}) + \lambda\|\widehat{\Theta}\|_* \leq L(\Theta^\star) + \lambda\|\Theta^\star\|_*.$$

Hence, by Hölder's inequality,

$$\begin{aligned} 0 \le D &= L(\widehat{\Theta}) - L(\Theta^\star) - \nabla L(\Theta^\star) \cdot \Delta \\ &\le ||\nabla L(\Theta^\star)||_2 \, ||\Delta||_* + \lambda \left(\|\Theta^\star\|_* - \|\widehat{\Theta}\|_* \right). \end{aligned} \tag{11}$$

Since $||\nabla L(\Theta^\star)||_2 \le \lambda$, triangle inequality in (11) yields

$$D \le ||\nabla L(\Theta^\star)||_2 \, ||\Delta||_* + \lambda ||\Delta||_* \le 2\lambda \, ||\Delta||_* .$$

Together with Lemma 3 and (10) (for the lower bound) and our choice of λ (for the upper bound), this yields

$$\frac{1}{8e^{4\gamma}\rho mnK} \, ||\Delta||_{\mathrm{F}}^2 \le D \le 16\sqrt{\tau} \sqrt{\frac{\rho K (m+n)\log(m+n)}{mnN}} \, ||\Delta||_* . \tag{12}$$

Hence recalling $\gamma = \frac{\alpha}{\sqrt{mn}}$,

$$||\Delta||_{\mathrm{F}}^2 \le 128\sqrt{\tau}\alpha e^{\frac{4\alpha}{\sqrt{mn}}} (\rho K)^{3/2} \sqrt{\frac{(m+n)\log(m+n)}{N}} \, ||\Delta||_* . \tag{13}$$

Let's return to (11). Since $||\nabla L(\Theta^\star)||_2 \le \lambda/2$, we have

$$\begin{aligned} 0 &\le \|\nabla L(\Theta^\star)\|_2 \|\Delta\|_* + \lambda \left(\|\Theta^\star\|_* - \|\widehat{\Theta}\|_* \right) \\ &\le \lambda \left(\frac{1}{2}\|\Delta\|_* + 2\|\overline{\Theta}_r^\star\|_* + \|\Delta'\|_* - \|\Delta''\|_* \right) \\ &\le \lambda \left(2\|\overline{\Theta}_r^\star\|_* + \frac{3}{2}\|\Delta'\|_* - \frac{1}{2}\|\Delta''\|_* \right), \end{aligned}$$

and so $||\Delta''||_* \le 3\,||\Delta'||_* + 4\|\overline{\Theta}_r^\star\|_*$. Let $\zeta = \|\overline{\Theta}_r^\star\|_* = \sum_{j=r+1}^{\min\{m,n\}} \sigma_j$. Since $||\Delta||_{\mathrm{F}} - ||\Delta'||_{\mathrm{F}} = ||\Gamma_{22}||_{\mathrm{F}} \ge 0$,

$$\begin{aligned} ||\Delta||_* &\le ||\Delta'||_* + ||\Delta''||_* = 4\,||\Delta'||_* + 4\zeta \\ &\le 8 \max\{||\Delta'||_*, \zeta\} \le 8 \max\left\{ \sqrt{2r}\,||\Delta'||_{\mathrm{F}}, \zeta \right\} \\ &\le 8 \max\left\{ \sqrt{2r}\,||\Delta||_{\mathrm{F}}, \zeta \right\} \\ &\le 16 \max\left\{ \sqrt{r}\,||\Delta||_{\mathrm{F}}, \zeta \right\}. \end{aligned} \tag{14}$$

If $\sqrt{r}\|\Delta\|_{\mathrm{F}} \ge \zeta$, substitute (14) in (13) to see

$$||\Delta||_{\mathrm{F}} \le 2048\sqrt{\tau}\alpha e^{\frac{4\alpha}{\sqrt{mn}}} K^{3/2} \sqrt{\frac{r(m+n)\log(m+n)}{N}}.$$

Otherwise (if $\sqrt{r}\|\Delta\|_{\mathrm{F}} < \zeta$), substitute (14) in (13) and take the square root to see

$$\begin{aligned} ||\Delta||_{\mathrm{F}} &\le \sqrt{2048\sqrt{\tau}\alpha e^{\frac{4\alpha}{\sqrt{mn}}} K^{3/2} \sqrt{\frac{\zeta(m+n)\log(m+n)}{N}}} \\ &\le 2048\sqrt{\tau}\alpha e^{\frac{4\alpha}{\sqrt{mn}}} K^{3/4} \left(\frac{\zeta d \log d}{N} \right)^{1/4}. \end{aligned}$$

Combining yields the statement.

Our last step is to investigate the case where $\Delta \notin \mathcal{A}^\star$, *i.e.*,

$$\|\Delta\|_* > \max\left\{(18\tau)^{1/4},\, 480\right\} \frac{1}{\sqrt{\rho^3 K m n \gamma}} \sqrt{\frac{N}{\sqrt{mn}\log(m+n)}} \|\Delta\|_{\mathrm{F}}^2.$$

Rewriting and introducing redundant terms greater than 1, we recover (13) and the proof follows as before from that point. $\square$

Proof of Lemma 2 Since S_t is symmetric, we have that

$$\mathbb{E}\left[Y_t(\Theta)\right] = \sum_{i=1}^m \mu_i \sum_{j=1}^n \frac{1}{n}\Theta_{ij}^2 \geq \frac{1}{\rho m n}\left|\left|\Theta\right|\right|_F^2.$$

Define

$$\tilde{\mathcal{M}}_{\Gamma,\nu} = \sup_{\Theta\in\mathcal{A}_{\Gamma,\nu}} \frac{1}{N}\sum_{t=1}^N \left(\mathbb{E}Y_t(\Theta) - Y_t(\Theta)\right)$$

and note that $\tilde{\mathcal{M}}_{\Gamma,\nu} \geq \mathcal{M}_{\Gamma,\nu}$. Then, letting $Y_t'(\Theta)$ be an identical and independent replicate of $Y_t(\Theta)$ and letting ϵ_t be iid Rademacher random variables independent of all else, we have

$$\begin{aligned}
\mathbb{E}\tilde{\mathcal{M}}_{\Gamma,\nu} &= \mathbb{E}\left[\sup_{\Theta\in\mathcal{A}_{\Gamma,\nu}} \frac{1}{N}\sum_{t=1}^N \left(\mathbb{E}Y_t'(\Theta) - Y_t(\Theta)\right)\right] \leq \mathbb{E}\left[\sup_{\Theta\in\mathcal{A}_{\Gamma,\nu}} \frac{1}{N}\sum_{t=1}^N \left(Y_t'(\Theta) - Y_t(\Theta)\right)\right] \\
&= \mathbb{E}\left[\sup_{\Theta\in\mathcal{A}_{\Gamma,\nu}} \frac{1}{N}\sum_{t=1}^N \epsilon_t\left(Y_t'(\Theta) - Y_t(\Theta)\right)\right] \leq 2\mathbb{E}\left[\sup_{\Theta\in\mathcal{A}_{\Gamma,\nu}} \frac{1}{N}\sum_{t=1}^N \epsilon_t Y_t(\Theta)\right] \\
&= 2\mathbb{E}\left[\sup_{\Theta\in\mathcal{A}_{\Gamma,\nu}} \frac{1}{N}\sum_{t=1}^N \epsilon_t \frac{1}{K_t}\left|\left|e_{i_t}^T\Theta\sum_{j\in S_t} e_j\right|\right|_2^2\right].
\end{aligned}$$

Note that

$$\sup_{v,v'\in[-\gamma,\gamma]^k\times\{0\}^{n-k}} \frac{||v||_2^2 - ||v'||_2^2}{||v-v'||_\infty} = 2\gamma k,$$

i.e., $||\cdot||_2^2$ is $2\gamma k$-Lipschitz with respect to the ∞-norm on a domain in $[-\gamma,\gamma]^n$ where only k entries are nonzero. Therefore, by Lemma 7 of Bertsimas and Kallus (2014) and by Hölder's inequality, letting $W_t = \sum_{j\in S_t}\epsilon_{tj}e_{i_t}e_j^T$ where ϵ_{tj} are new iid Rademacher random variables independent of all else,

$$\mathbb{E}\tilde{\mathcal{M}}_{\Gamma,\nu} \leq 4\gamma\mathbb{E}\left[\sup_{\Theta\in\mathcal{A}_{\Gamma,\nu}} \frac{1}{N}\sum_{t=1}^N W_t\cdot\Theta\right] \leq 4\gamma\mathbb{E}\left|\left|\frac{1}{N}\sum_{t=1}^N W_t\right|\right|_2 \sup_{\Theta\in\mathcal{A}_{\Gamma,\nu}}||\Theta||_*.$$

Note that

$$\begin{aligned}
&||W_t||_2 \leq \sqrt{K} \\
&\mathbb{E}\left[W_tW_t^T\middle|S_t,i_t\right] = K_t e_{i_t}e_{i_t}^T \qquad \text{and hence} \qquad \left|\left|\mathbb{E}\left[W_tW_t^T\right]\right|\right|_2 \leq \rho K/m \\
&\mathbb{E}\left[W_t^TW_t\middle|S_t,j_t\right] = \sum_{j\in S_t} e_je_j^T \qquad \text{and hence} \qquad \left|\left|\mathbb{E}\left[W_t^TW_t\right]\right|\right|_2 \leq K/n
\end{aligned}$$

Hence, the matrix Bernstein inequality (Theorem 1.6 of Tropp (2012)) gives that $\left|\left|\frac{1}{N}\sum_{t=1}^N W_t\right|\right|_2 \geq \delta$ with probability at most

$$(m+n)\max\left\{e^{-\frac{N\delta^2\min\{m,n\}}{4\rho K}}, e^{-\frac{\delta}{\sqrt{K}}}\right\}.$$

Setting the probability to $1/\sqrt{mn\min\{m,n\}}$ and using $N \leq mn\log(m+n)$,

$$\begin{aligned}\mathbb{E}\left[\left|\left|\frac{1}{N}\sum_{t=1}^N W_t\right|\right|_2\right] &\leq \sqrt{\frac{K}{mn\min\{m,n\}}} + 2\sqrt{\frac{\rho K\left(\frac{1}{2}\log(m)+\frac{1}{2}\log(n)\frac{1}{2}\log(\min\{m,n\})+\log(m+n)\right)}{N\min\{m,n\}}} \\ &\leq \sqrt{\frac{K\log(m+n)}{N\min\{m,n\}}} + 2\sqrt{\frac{\frac{5}{2}\rho K\log(m+n)}{N\min\{m,n\}}} \leq 5\sqrt{\frac{\rho K\log(m+n)}{N\min\{m,n\}}}\end{aligned}$$

Putting it all together, we get,

$$\mathbb{E}\tilde{\mathcal{M}}_{\Gamma,\nu} \leq 20\gamma\sqrt{\frac{\rho K\log(m+n)}{N\min\{m,n\}}}\frac{\nu}{60\sqrt{\rho^3 K mn\gamma}}\sqrt{\frac{N}{(m+n)\log(m+n)}}\Gamma^2 \leq \frac{\nu}{3}\frac{\Gamma^2}{\rho mn}.$$

Next we use this to prove the concentration of $\tilde{\mathcal{M}}_{\Gamma,\nu}$. Let $\tilde{\mathcal{M}}'_{\Gamma,\nu}$ be a replicate of $\tilde{\mathcal{M}}_{\Gamma,\nu}$ with $i'_t = i_t$, $S'_t = S_t$ for all t except t'. Then the difference $\tilde{\mathcal{M}}_{\Gamma,\nu} - \tilde{\mathcal{M}}'_{\Gamma,\nu}$ is bounded by

$$\frac{1}{N}\sup_{\Theta\in\mathcal{A}_{\Gamma,\nu}}\left(\frac{1}{K_{t'}}\sum_{j\in S_{t'}}\Theta^2_{i_{t'}j} - \frac{1}{K_{t'}}\sum_{j\in S'_{t'}}\Theta^2_{i'_{t'}j}\right) \leq \frac{1}{N}\left(\gamma^2 - 0\right) = \frac{\gamma^2}{N}.$$

Hence, by McDiarmid's inequality, we have

$$\mathbb{P}\left(\tilde{\mathcal{M}}_{\Gamma,\nu} \geq \nu\frac{\Gamma^2}{\rho mn}\right) \leq \mathbb{P}\left(\tilde{\mathcal{M}}_{\Gamma,\nu} - \mathbb{E}\mathcal{M}_{\Gamma,\nu} \geq \frac{2\nu}{3}\frac{\Gamma^2}{\rho mn}\right) \leq \exp\left(-\frac{8}{9}\frac{N\nu^2\Gamma^4}{\gamma^4\rho^2m^2n^2}\right),$$

which yields the result, given $\tilde{\mathcal{M}} \geq \mathcal{M}$. □

Proof of Lemma 3 Let $\tau' = K^2\rho^4\max\{\tau, 480^4/18\}$. Since $\|\cdot\|_* \geq \|\cdot\|_{\mathrm{F}}$, we have $\inf_{\Theta\in\mathcal{A}^\star}\|\Theta\|_{\mathrm{F}} \geq \eta := (18\tau')^{1/4}\sqrt{\rho mn}\gamma\sqrt{\sqrt{mn}\log(m+n)/N}$. Set $\nu = 60(18\tau')^{-1/4}$ and $\beta = \sqrt{1/(2\nu)}$. Since $\tau' \geq 480^4/18$ we have $\nu \leq 1/8$ and $\beta \geq 2 > 1$. Let $\mathcal{A}_l = \mathcal{A}^\star \cap \{\eta\beta^{l-1} \leq ||\Theta||_{\mathrm{F}} \leq \eta\beta^l\}$ and note that $\mathcal{A}^\star = \bigcup_{l=1}^\infty A_l$ and $\mathcal{A}_l \subset \mathcal{A}_{\beta^l\eta,\nu}$.

If $\Theta \in \mathcal{A}_l$ has $\frac{1}{N}\sum_{t=1}^N Y_t(\Theta) < \frac{1}{2\rho mn}||\Theta||^2_{\mathrm{F}}$ then

$$\frac{1}{\rho mn}||\Theta||^2_{\mathrm{F}} - \frac{1}{N}\sum_{t=1}^N Y_t(\Theta) > \frac{1}{2\rho mn}||\Theta||^2_{\mathrm{F}} \geq \frac{1}{2\rho mn}\left(\beta^{l-1}\eta\right)^2 = \frac{1}{\rho mn}\frac{1}{2\beta^2}(\beta^l\eta)^2 = \frac{\nu}{\rho mn}(\beta^l\eta)^2.$$

Therefore, the probability that the event in the statement of theorem is invalid is bounded by

$$\begin{aligned}&\min\left\{1, \sum_{l=1}^\infty \mathbb{P}\left(\mathcal{M}_{\beta^l\eta,\nu} > \frac{\nu}{\rho mn}\left(\beta^\ell\eta\right)^2\right)\right\} \\ &\leq \min\left\{1, \sum_{l=1}^\infty \exp\left(-\frac{8}{9}\frac{\nu^2\beta^{4\ell}\eta^4 N}{\rho^2m^2n^2\gamma^4}\right)\right\}\end{aligned}$$

$$
\begin{aligned}
&\leq \min\left\{1, \sum_{l=1}^{\infty} \exp\left(-\frac{1}{72}\frac{4^{\ell}\eta^4 N}{\rho^2 m^2 n^2 \gamma^4}\right)\right\} \\
&\leq \min\left\{1, \sum_{l=1}^{\infty} \exp\left(-\frac{1}{72}\frac{4^{l}\eta^4 N}{\rho^2 m^2 n^2 \gamma^4}\right)\right\} \\
&\leq \min\left\{1, \sum_{l=1}^{\infty} \exp\left(-\frac{1}{18}\frac{\eta^4 N}{\rho^2 m^2 n^2 \gamma^4} l\right)\right\} \\
&= \min\left\{1, \left(\exp\left(\frac{1}{18}\frac{\eta^4 N}{\rho^2 m^2 n^2 \gamma^4}\right) - 1\right)^{-1}\right\} \\
&\leq 2\exp\left(-\frac{1}{18}\frac{\eta^4 N}{\rho^2 m^2 n^2 \gamma^4}\right) \\
&= 2\exp\left(-\tau' mn(\log(m+n))^2/N\right) \\
&\leq 2(m+n)^{-\tau'} \leq 2(m+n)^{-\tau},
\end{aligned}
$$

using Lemma 2 and $N \leq mn\log(m+n)$. $\square$

Proof of Lemma 4 Let $\psi_{tj} = e^{X_{tj}\cdot\Theta^\star}$ and $G_t = \frac{\sum_{j\in S_t}\psi_{tj}X_{tj}}{1+\sum_{j\in S_t}\psi_{tj}} - X_{tj_t}$ so that $\nabla L(\Theta^\star) = \frac{1}{N}\sum_{t=1}^{N} G_t$. Note that because j_t is drawn according to $\Theta^\star$ and $X_{t0} = 0$, we have that $\mathbb{E}\left[G_t \middle| i_t, S_t\right] = 0$ and hence $\mathbb{E}G_t = 0$. Note that $\|G_t\|_2 \leq \sqrt{2}$. Moreover,

$$
G_t G_t^T = e_{i_t} e_{i_t}^T \left(1 - \frac{2\psi_{tj_t}}{1+\sum_{j\in S_t}\psi_{tj}} + \frac{\sum_{j\in S_t}\psi_{tj}^2}{\left(1+\sum_{j\in S_t}\psi_{tj}\right)^2}\right).
$$

Since by Jensen's inequality the multiplier in the parentheses is no greater than 2, we get $\left|\left|\mathbb{E}\left[G_t G_t^T\right]\right|\right|_2 \leq \frac{2\rho}{m} \leq \frac{2\rho K}{m}$. Letting $y_{tj} = \mathbb{I}[j = j_t]$, we have

$$
G_t^T G_t = \sum_{j\in S_t k\in S_t} e_j e_k^T \left(y_{tj}y_{tk} - \frac{2y_{tj}\psi_{tk}}{1+\sum_{l\in S_t}\psi_{tl}} + \frac{\psi_{tj}\psi_{tk}}{\left(1+\sum_{l\in S_t}\psi_{tl}\right)^2}\right).
$$

Noting that $y_{tj} \geq 0$, $\psi_{tj} \geq 0$, and $y_{tj}y_{tk} \leq \mathbb{I}[j=k]$, we see that

$$
\left|\left|\mathbb{E}\left[G_t^T G_t\right]\right|\right|_2 \leq \left|\left|\mathbb{E}\left[\sum_{j\in S_t} e_j e_j^T\right]\right|\right|_2 + \left|\left|\mathbb{E}\left[\frac{1}{K_t^2}\left(\sum_{j\in S_t} e_j\right)\left(\sum_{j\in S_t} e_j\right)^T\right]\right|\right|_2 \leq \frac{K}{n} + \frac{1}{n} \leq \frac{2\rho K}{n}.
$$

Hence, by the matrix Bernstein inequality (Theorem 1.6 of Tropp (2012)),

$$
\left|\left|\frac{1}{N}\sum_{t=1}^{N} G_t\right|\right|_2 \leq 2\sqrt{2(\tau+1)}\sqrt{\frac{\rho K\log(m+n)}{\min\{m,n\}N}} \leq 4\sqrt{\tau}\sqrt{\frac{\rho K(m+n)\log(m+n)}{mnN}},
$$

with probability at least $1-(m+n)^{-\tau}$. $\square$

This lemma concludes the proof of Theorem 3. We now prove that the estimate $\hat{mu}$ for the customer type distribution is consistent.

Proof of Theorem 4 We prove this by cases, depending on which term in the min is smaller. Note that this is not a random event so we can choose which bound to use a priori, yielding the min.

First we deal with the second term. By union bound and Hoeffding's inequality,

$$\mathbb{P}\left(\left|\left|\hat{\mu}-\mu^\star\right|\right|_q > \eta\right) \le \sum_{i=1}^m \mathbb{P}\left(|\hat{\mu}_i - \mu_i^\star| > \eta/m^{1/q}\right) \le 2m\exp(-2n\eta^2/m^{2/q})$$

Next we deal with the first term. Note that $\left|\left|\cdot\right|\right|_q \le \left|\left|\cdot\right|\right|_1$ so it is sufficient to prove this for $q=1$. Let $I_t = e_{i_t}$ be the type indicator observation and let I be a generic draw of I_t from μ. Then $\hat{\mu} = \frac{1}{N}\sum_{t=1}^N I_t$. By Bartlett and Mendelson (2003), we have that with probability at least $1-\nu$,

$$\left|\left|\mu-\hat{\mu}\right|\right|_1 = \sup_{||v||_\infty \le 1}\left(\mathbb{E}\left[v^T I\right] - \frac{1}{N}\sum_{t=1}^N v^T I_t\right) \le 2\widehat{\mathfrak{R}}_N + \sqrt{\frac{-\log\nu}{2N}},$$

where $\widehat{\mathfrak{R}}_N$ is the empirical Rademacher complexity

$$\widehat{\mathfrak{R}}_N = \frac{1}{2^N}\sum_{\epsilon\in\{-1,+1\}^N}\sup_{||v||_\infty\le 1}\frac{1}{N}\sum_{t=1}^N \epsilon_t v^T I_t.$$

By linearity and duality of norms, we have

$$\begin{aligned}
\widehat{\mathfrak{R}}_N &= \frac{1}{2^N}\sum_{\epsilon\in\{-1,+1\}^N}\sup_{||v||_\infty\le 1} v^T\left(\frac{1}{N}\sum_{t=1}^N \epsilon_t I_t\right)\\
&= \frac{1}{2^N}\sum_{\epsilon\in\{-1,+1\}^N}\left|\left|\frac{1}{N}\sum_{t=1}^N \epsilon_t I_t\right|\right|_1\\
&= \frac{1}{N}\sum_{i=1}^m \frac{1}{2^N}\sum_{\epsilon\in\{-1,+1\}^N}\left|\sum_{t=1}^N \mathbb{I}\left[i_t=i\right]\epsilon_t\right|\\
&= \frac{1}{N}\sum_{i=1}^m \frac{1}{2^{N\hat{\mu}_i}}\sum_{\epsilon\in\{-1,+1\}^{N\hat{\mu}_i}}\left|\sum_{t=1}^{N\hat{\mu}_i}\epsilon_t\right|\\
&= \frac{1}{N}\sum_{i=1}^m \frac{1}{2^{N\hat{\mu}_i-1}}\left\lceil\frac{N\hat{\mu}_i}{2}\right\rceil\binom{m}{\lceil\frac{N\hat{\mu}_i}{2}\rceil}\\
&\le \frac{1}{N}\sum_{i=1}^m\sqrt{N\hat{\mu}_i} = \frac{1}{\sqrt{N}}\sum_{i=1}^m\sqrt{\hat{\mu}_i} \le \sqrt{\frac{m}{N}}.
\end{aligned}$$

Therefore, using the concavity of square root, we have that

$$\mathbb{P}\left(\left|\left|\mu-\hat{\mu}\right|\right|_1 > \eta\right) \le e^{m-N\eta^2/64}.$$

Rearranging yields the result. □

8. Proofs Omitted from Section 2.3

Proof of Lemma 1 Given Θ feasible in (6) with $\operatorname{rank}(\Theta) \le \tilde{r}$, write its SVD $\Theta = \tilde{U}\Sigma\tilde{V}^T$, where $\Sigma \in \mathbb{R}^{\tilde{r}\times\tilde{r}}$ is diagonal and U, V unitary. Letting $U = \tilde{U}\Sigma^{1/2}$ and $V = \tilde{V}\Sigma^{1/2}$, we obtain a feasible solution to (7) with the same objective value Θ has in (6). Conversely, given U, V feasible in (7), let $\Theta = UV^T$. Note $\operatorname{rank}(\Theta) \le \tilde{r}$ and

$$\begin{aligned}\|\Theta\|_* &\le \|UV^T\|_* = \operatorname{tr}(\Sigma) = \operatorname{tr}(\tilde{U}^T UV^T\tilde{V}) \\ &\le \|\tilde{U}^T U\|_{\mathrm{F}}\|V^T\tilde{V}\|_{\mathrm{F}} \le \|U\|_{\mathrm{F}}\|V\|_{\mathrm{F}} \le \frac{1}{2}\|U\|_{\mathrm{F}}^2 + \frac{1}{2}\|V\|_{\mathrm{F}}^2.\end{aligned}$$

Hence, Θ has objective value no worse than (U, V). □

9. Proofs of Regret Bounds

Before presenting the proofs that a structure aware algorithm can achieve regret sublinear in the number of matrix entries mn, we formally prove that a structure ignorant algorithm must have regret linear in mn. This result, Theorem 5, makes use of Theorem 1 in Sauré and Zeevi (2013), which we restate here in our notation:

THEOREM 8 **(Sauré and Zeevi (2013), Theorem 1)**. *Consider a LRMMNL with one type, $m = 1$. Under any consistent policy π for choosing the sets S_t,*

$$\mathbb{E}\operatorname{Regret}_T(\pi) \ge C_1 + C_2\tilde{\mathcal{N}}\log T$$

for all T, where

$$\tilde{\mathcal{N}} = \left|\left\{j \notin S^\star(W_i, \Theta_i^\star; K) :\ \exists \theta \in \mathbb{R}^n,\, j \in S^\star(W_i, \theta; K),\, \theta_{j'} = \Theta_{ij}^\star\ \forall j' \in S^\star(W_i, \Theta_i^\star; K)\right\}\right|$$

is the set of potentially optimal items, and C_1 and C_2 are constants independent of n and T.

Note that Sauré and Zeevi (2013) prove this lower bound only for policies that do not depend on the time horizon T, since their proof requires taking $T \to \infty$. However, this proof is easily extended to a lower bound even for policies that depend on the time horizon by considering a sequence of policies for each T.

Proof of Theorem 5 First, use a Chernoff bound to see that the number of times T_i that type i is chosen is larger than $\frac{T\mu_i}{2}$ with high probability:

$$\mathbb{P}\left(T_i > \frac{T\mu_i}{2}\right) \ge 1 - \exp\left(\frac{-1}{8}T\mu_i\right).$$

Recall that $T = \Omega(m^{1+\epsilon})$ grows superlinearly in the number of types m. We use a union bound to show that for large T, with high probability we have $T_i \geq \frac{T\mu_i}{2}$ for every i:

$$\begin{aligned}
\mathbb{P}\left(T_i \geq \frac{T\mu_i}{2} \forall i\right) &\geq 1 - \sum_{i=1}^{m} \exp\left(\frac{-1}{8} T\mu_i\right) \\
&\geq 1 - \sum_{i=1}^{m} \exp\left(\frac{-T}{8\rho m}\right) \\
&\geq 1 - \sum_{i=1}^{m} \exp\left(\frac{-m^{\epsilon}}{8\rho}\right) \\
&\geq 1 - m \exp\left(\frac{-m^{\epsilon}}{8\rho}\right),
\end{aligned}$$

which converges to 1 as m increases. In particular, for sufficiently large T and m, $\mathbb{P}\left(T_i \geq \frac{T\mu_i}{2} \forall i\right) > \frac{1}{2}$. Let's suppose below that T and m are large enough that this bound holds.

Then conditioning on the event that $T_i \geq \frac{T\mu_i}{2}$ for every i,

$$\begin{aligned}
\mathbb{E}\left[\sum_{t=1}^{T} r_t\right] &\geq \mathbb{E}\left[\sum_{t=1}^{T} r_t \,|\, T_i \geq \frac{T\mu_i}{2} \ \forall i\right] \mathbb{P}\left(T_i \geq \frac{T\mu_i}{2} \ \forall i\right) \\
&\geq \frac{1}{2}\mathbb{E}\left[\sum_{t=1}^{T} r_t \,|\, T_i \geq \frac{T\mu_i}{2} \ \forall i\right] \\
&= \frac{1}{2}\mathbb{E}\left[\sum_{i=1}^{m} \sum_{t:\ i_t = i} r_t \,|\, T_i \geq \frac{T\mu_i}{2} \ \forall i\right] \\
&= \frac{1}{2}\left(\sum_{i=1}^{m} \mathbb{E}\left[\sum_{t:\ i_t = i} r_t \,|\, T_i \geq \frac{T\mu_i}{2} \ \forall i\right]\right) \\
&= \frac{1}{2}\left(\sum_{i=1}^{m} \mathbb{E}\left[\sum_{t:\ i_t = i} r_t \,|\, T_i \geq \frac{T\mu_i}{2}\right]\right)
\end{aligned}$$

Now, using Theorem 8 and the assumption that $\tilde{\mathcal{N}} \geq \nu n$,

$$\begin{aligned}
\mathbb{E}\left[\sum_{t=1}^{T} r_t\right] &\geq \frac{1}{2}\left(\sum_{i=1}^{m} C_1 + C_2(\nu n)\log(T\mu_i/2)\right) \\
&\geq \frac{1}{2}\left(\sum_{i=1}^{m} C_1 + C_2(\nu n)\log\left(\frac{T}{2\rho m}\right)\right) \\
&\geq \frac{1}{2}\left(C_1 m + C_2'(\nu n)\sum_{i=1}^{m}(\log T - \log(2\rho m))\right).
\end{aligned}$$

Hence if $m = o(T)$, we have

$$\mathbb{E}\left[\sum_{t=1}^{T} r_t\right] = \Omega(mn \log T).$$

(Otherwise, we have $\mathbb{E}[\sum_{t=1}^{T} r_t] = \Omega(mn\log(T/m))$.) $\square$

Now we show that our structure aware algorithms produce regret sublinear in mn.

Proof of Theorem 6 Note that, for any S and $j \in S$,

$$0 \leq \frac{\partial p_j(S;\theta)}{\partial \theta_j} = \frac{e^{\theta_j} \sum_{j' \in S,\, j' \neq j} e^{\theta_{j'}}}{\left(\sum_{j' \in S} e^{\theta_{j'}}\right)^2} \leq \frac{|S|-1}{|S|^2} \leq \frac{1}{4},$$

and that, moreover, for any $j' \in S$, $j' \neq j$,

$$0 \geq \frac{\partial p_j(S;\theta)}{\partial \theta_{j'}} = \frac{-e^{\theta_j+\theta_{j'}}}{\left(\sum_{j'' \in S} e^{\theta_{j''}}\right)^2} \geq \frac{-e^{\theta_j+\theta_{j'}}}{\left(e^{\theta_j}+e^{\theta_{j'}}\right)^2} \geq -\frac{1}{4},$$

whereas, for $j' \notin S$, clearly $\frac{\partial p_j(S;\theta)}{\partial \theta_{j'}} = 0$. Therefore, $||\nabla_\theta p_j(S;\theta)||_2 \leq \frac{1}{4}\sqrt{|S|}$ for any $\theta \in \mathbb{R}^n$. This means that for any S, $j \in S$, θ, and θ',

$$|p_j(S;\theta) - p_j(S;\theta')| \leq \frac{1}{4}\sqrt{|S|}\,||\theta-\theta'||_2\,.$$

It follows that for any S with $|S| \leq K$, w, θ, and θ',

$$|F(S;w,\theta) - F(S;w,\theta')| \leq \sum_{j \in S} |w_j|\,|p_j(S;\theta) - p_j(S;\theta')| \leq \frac{1}{4}\,||w||_\infty K^{3/2}\,||\theta-\theta'||_2\,.$$

Therefore, letting $S \in S^\star(w,\theta;K)$ and $S' \in S^\star(w,\theta';K)$,

$$\begin{aligned} F(S';w,\theta) &\geq F(S';w,\theta') - \frac{1}{4}\,||w||_\infty K^{3/2}\,||\theta-\theta'||_2 \\ &\geq F(S;w,\theta') - \frac{1}{4}\,||w||_\infty K^{3/2}\,||\theta-\theta'||_2 \\ &\geq F(S;w,\theta) - \frac{1}{2}\,||w||_\infty K^{3/2}\,||\theta-\theta'||_2\,. \end{aligned}$$

It follows that

$$\left(||\theta-\theta'||_2 \leq 2\delta(w,\theta;K)K^{-3/2}/\,||w||_\infty\right) \implies \left(S^\star(w,\theta';K) = S^\star(w,\theta;K)\right).$$

Noting that $||W_i||_\infty \leq ||W||_\infty$, $\left|\left|\widehat{\Theta}_i - \Theta_i^\star\right|\right|_2 \leq \left|\left|\widehat{\Theta} - \Theta^\star\right|\right|_{\mathrm{F}}$, $0 < \delta \leq \delta(W_i,\Theta_i^\star;K)$, $T - T_0 \leq T$, and that the per-step regret is at most $||W||_\infty$, we have that

$$\mathrm{Regret}_{\Theta^\star}^{\pi_1(T_0,\lambda)} \leq ||W||_\infty T_0 + ||W||_\infty T \times \mathbb{P}\left(\left|\left|\widehat{\Theta} - \Theta^\star\right|\right|_{\mathrm{F}} > 2\delta K^{-3/2}/\,||W||_\infty\right).$$

Applying Theorem 3 with $\tau = \log(T)/\log(m+n)$ and using $\alpha \geq \alpha/\sqrt{mn}$ gives the result. □

Proof of Theorem 7 Note that if $j \in S$ then

$$\left|\frac{\partial F(S;W,\Theta,\mu)}{\partial \Theta_{ij}}\right| \leq \mu_i \sum_{j' \in S} \frac{w_{ij'}}{4} \leq \frac{K\rho\,||W||_\infty}{4m},$$

and if $j \notin S$ then $\frac{\partial F(S;W,\Theta,\mu)}{\partial \Theta_{ij}} = 0$. Moreover,

$$\left|\frac{\partial F(S;W,\Theta,\mu)}{\partial \mu_i}\right| = \left|\frac{\sum_{j\in S} e^{\Theta_{ij}} w_{ij}}{1+\sum_{j\in S} e^{\Theta_{ij}}}\right| \le ||W||_\infty .$$

It follows that for any S with $|S| \le K$, W, Θ, μ, Θ', and μ',

$$\begin{aligned}|F(S;W,\Theta,\mu) - F(S;W,\Theta',\mu')| &\le \frac{K\rho\,||W||_\infty}{4m}\sum_{i=1}^m \sum_{j\in S} \left|\Theta_{ij} - \Theta'_{ij}\right| + ||W||_\infty \sum_{i=1}^m |\mu - \mu'| \\ &\le \frac{1}{4}K^{3/2}\rho\,||W||_\infty \,||\Theta - \Theta'||_{\mathrm{F}} + ||W||_\infty \,||\mu - \mu'||_1 .\end{aligned}$$

Therefore, since the per-step regret is at most $||W||_\infty$, we have that

$$\begin{aligned}\mathrm{Regret}_{\Theta^\star}^{\pi_2(T_0,\lambda)} &\le ||W||_\infty T_0 + ||W||_\infty T \times \mathbb{P}\left(\frac{1}{4}K^{3/2}\rho\,||W||_\infty \left|\left|\widehat{\Theta} - \Theta^\star\right|\right|_{\mathrm{F}} + ||W||_\infty \,||\hat{\mu} - \mu||_1 > \delta\right) \\ &\le \omega T_0 + \omega T \times \left(\mathbb{P}\left(||\Theta - \Theta'||_{\mathrm{F}} > \frac{2\delta}{K^{3/2}\rho\omega}\right) + \mathbb{P}\left(||\mu - \mu'||_1 > \frac{\delta}{2\omega}\right)\right).\end{aligned}$$

Applying Theorem 3 with $\tau = \log(T)/\log(m+n)$ we get that $\mathbb{P}\left(||\Theta - \Theta'||_{\mathrm{F}} > \frac{4\beta\delta}{K^{3/2}\rho\omega}\right) \le 3/T$ given our choice of T_0.

By Theorem 4, $\mathbb{P}\left(||\mu - \hat{\mu}||_1 > \eta\right) \le e^{m - T_0\eta^2/64}$. Since $T_0 \ge \frac{256\omega(m+\log(T))}{\delta^2}$ we have $\mathbb{P}\left(||\mu - \mu'||_1 > \frac{\delta}{2\omega}\right) \le 1/T$. Finally, $\alpha \ge \alpha/\sqrt{mn}$ yields the result. $\square$